

Multi-Site Network and Security Services with NSX-T

Implement Network Security,
Stateful Services, and Operations

Iwan Hoogendoorn

Apress®

WOW! eBook
www.wowebook.org

Multi-Site Network and Security Services with NSX-T

**Implement Network Security,
Stateful Services,
and Operations**

Iwan Hoogendoorn

Apress®

Multi-Site Network and Security Services with NSX-T: Implement Network Security, Stateful Services, and Operations

Iwan Hoogendoorn
ROTTERDAM, Zuid-Holland, The Netherlands

ISBN-13 (pbk): 978-1-4842-7082-0
<https://doi.org/10.1007/978-1-4842-7083-7>

ISBN-13 (electronic): 978-1-4842-7083-7

Copyright © 2021 by Iwan Hoogendoorn

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr
Acquisitions Editor: Aditee Mirashi
Development Editor: Laura Berendson
Coordinating Editor: Aditee Mirashi

Cover designed by eStudioCalamar

Cover image designed by Freepik (www.freepik.com)

Distributed to the book trade worldwide by Springer Science+Business Media New York, 1 New York Plaza, Suite 4600, New York, NY 10004-1562, USA. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a Delaware corporation.

For information on translations, please e-mail booktranslations@springernature.com; for reprint, paperback, or audio rights, please e-mail bookpermissions@springernature.com.

Apress titles may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Print and eBook Bulk Sales web page at <http://www.apress.com/bulk-sales>.

Any source code or other supplementary material referenced by the author in this book is available to readers on GitHub via the book's product page, located at www.apress.com/978-1-4842-7082-0. For more detailed information, please visit <http://www.apress.com/source-code>.

Printed on acid-free paper

*The world is full of fathers, but there are very few men
worthy of being called daddy.*

*Becoming a father is the most beautiful thing that happened
to me in life. Zaara is an identical copy when it comes to my
personality. I am having arguments with “myself” multiple
times a day; I am frustrated by the things she does and says
and, in reality, I am getting frustrated with myself.
This brings up the question, how can people live with me?
Am I really like this?*

All credits go to my wife and daughter, I guess...

*To all the fathers in the world: Be there, be an example, and
be proud, but most importantly, make them proud of you!
Set an example so they can survive and find their way in
this “broken” world we live in. They are the future.*

*Zaara Ameerah Hoogendoorn, I am proud of you and I am
doing everything I can to make you proud of me as well.*

Table of Contents

About the Author	xvii
About the Technical Reviewers	xix
Acknowledgments	xxi
Introduction	xxiii
Chapter 1: NSX-T Security and Firewalls.....	1
Traditional Data Center Security	1
Data Center Security Requirements.....	2
NSX-T Micro-Segmentation	3
Zero Trust Security Model	4
NSX-T Micro-Segmentation Use Cases.....	5
NSX-T Micro-Segmentation Benefits	5
Enforcing the Zero Trust Security Model Using Micro-Segmentation	6
NSX-T Distributed Firewall.....	7
NSX-T Distributed Firewall (DFW) Features.....	8
NSX-T DFW Concepts	9
Security Policy Overview	9
DFW Security Policy.....	10
DFW Configuration.....	12
DFW Policy Advanced Setting Configuration	15
DFW Time-Based Security Policy Configuration	16

TABLE OF CONTENTS

DFW Rule Configuration.....	18
DFW Rule Parameters Configuration	19
DFW Source and Destination Definitions	20
Groups	20
Creating a Group and Adding Members.....	21
Service Specification Inside a DFW Rule.....	23
Add a Context Profile to a DFW Rule.....	27
Context Attributes Configuration	29
Scope of the DFW Rule.....	32
DFW Rule Actions	33
DFW Rule Settings.....	34
Enable/Disable a DFW Rule	35
Save and View DFW Rule Configurations.....	35
Roll Back Saved DFW Rule Configurations	37
Distributed Firewall (DFW) Architecture	41
DFW Architecture on the ESXi Hypervisor	42
NSX-T Gateway Firewall	44
Predefined NSX-T Gateway Firewall Categories.....	46
NSX-T Gateway Firewall Policy.....	47
NSX-T Gateway Firewall Policy Setting Configuration.....	48
NSX-T Gateway Firewall Rule Configuration	49
NSX-T Gateway Firewall Rule Setting Configuration	50
NSX-T Gateway Firewall Architecture.....	50
Summary.....	51
Chapter 2: NSX-T Advanced Security	53
Distributed Intrusion Detection	53
Distributed Intrusion Detection Use Case	54
Distributed Intrusion Detection Requirements.....	54

IDS Signatures54

IDS Profiles55

IDS Policy and Rules55

Distributed Intrusion Detection Architecture.....56

Distributed Intrusion Detection Configuration.....57

IDS Profile Configuration.....58

IDS Rules Configuration58

IDS Event Monitoring59

URL Analysis60

URL Analysis Use Case.....61

URL Analysis Requirements61

URL Categories.....61

Reputation Score and Severity.....61

URL Analysis Use Case.....62

Webroot as the Cloud Service63

URL Analysis Architecture63

URL Analysis Configuration65

URL Analysis Context Profiles65

Layer-7 Rule for DNS Traffic.....66

URL Analysis Dashboard67

Summary.....68

Chapter 3: NSX-T Service Insertion69

 Service Insertion69

 Endpoint Protection Use Cases.....70

 Network Introspection.....70

 North-South Network Introspection71

 East-West Network Introspection71

TABLE OF CONTENTS

Network Introspection Configuration	73
Service Registration.....	74
Service Deployment.....	75
Service Consumption	76
Service Profile Creation	77
Service Chain Creation	77
Redirection Policy Configuration	78
Endpoint Protection Overview.....	79
Endpoint Protection Use Cases	79
Endpoint Protection Process	79
New VMs: Automated Policy Enforcement	80
Virus and Malware: Automated Quarantine with Security Tags	81
Service Profile Creation for Endpoint Protection.....	82
Endpoint Protection Rules Configuration	83
Guest Introspection Architecture.....	83
Summary.....	86
Chapter 4: NSX-T NAT, DHCP, and DNS Services.....	87
Network Address Translation (NAT) Support	87
Source Network Address Translation (SNAT)	88
Destination Network Address Translation (DNAT)	89
Reflexive Network Address Translation (Reflexive NAT)	91
SNAT and DNAT Configuration	92
No SNAT or DNAT Rule Configuration	95
Reflexive NAT Configuration	95
NAT Packet Flow.....	96

NAT64.....	101
NAT64 Use Cases.....	103
NAT64 Requirements.....	103
NAT64 Limitations.....	104
NAT64 Tables	105
NAT64 Packet Flow.....	105
NAT64 Rule Configuration.....	108
DHCP Services	108
DHCP Architecture	109
DHCP Use Cases.....	110
DHCP Workflow.....	111
DHCP Server Profile Creation	111
DHCP Server on a Tier-0/Tier-1 Gateway.....	112
DHCP Configuration on a Segment.....	114
DHCP Server Status Verification	116
DHCP Relay Profile Creation	116
DHCP Relay Server on a Tier-0/Tier-1 Gateway Configuration.....	118
DNS Services	118
DNS Forwarder	119
DNS Forwarder Benefits	121
DNS Services and DNS Zones Configuration	121
DNS Forwarder Verification	123
Summary.....	124
Chapter 5: NSX-T Load Balancing.....	125
NSX-T Load Balancing Use Cases.....	125
Layer-4 Load Balancing	126
Layer-7 Load Balancing	126

TABLE OF CONTENTS

Load Balancer Component Relation.....	127
Load Balancer Architecture.....	128
Load Balancer Attachment to a Tier-1 Gateway.....	129
Virtual Servers	130
Profiles.....	130
Server Pools.....	131
Monitors.....	132
Load Balancer Deployment Modes	133
Inline Mode.....	135
One-Arm Mode	135
Load Balancing Configuration Steps.....	137
Load Balancer Creation.....	138
Virtual Server Creation.....	139
Layer-4 Virtual Server Creation	140
Layer-7 Virtual Server Creation.....	142
Application Profile Configuration	144
Persistence Profile Configuration.....	145
Layer-7 Load Balancer SSL Modes	146
Layer-7 SSL Profile Configuration	148
Layer-7 SSL Load Balancer Rules Configuration.....	149
Server Pool Creation	151
Load-Balancing Algorithms.....	152
SNAT Translation Mode Configuration.....	153
Active Monitoring Configuration.....	154
Passive Monitoring Configuration	155
Summary.....	156

Chapter 6: NSX-T VPN	157
NSX-T VPN Services.....	157
IPsec VPN Use Cases.....	157
IPsec VPN Methods.....	158
IPsec VPN Modes.....	160
IPsec Protocols and Algorithms.....	160
IPsec VPN Certificate-Based Authentication.....	161
IPsec VPN Dead Peer-to-Peer Detection.....	162
IPsec VPN Types.....	162
IPsec VPN Deployment Considerations.....	163
IPsec High Availability (HA).....	164
IPsec VPN Scalability.....	165
IPsec VPN Configuration.....	166
IPsec VPN Service Configuration.....	167
DPD Profile Configuration.....	168
IKE Profile Configuration.....	170
IPsec Profile Configuration.....	172
IPsec VPN Session Configuration.....	174
L2VPN.....	178
L2VPN Use Cases.....	178
NSX-T L2VPN.....	178
NSX-T L2VPN Deployment Considerations.....	178
NSX-T L2VPN Hub-and-Spoke.....	179
NSX-T L2VPN Packet Format.....	180
NSX-T L2VPN Packet Flow.....	180
Layer-2 VPN Scalability.....	181
NSX-T L2VPN Configuration Steps.....	182
IPsec Local Endpoint Configuration.....	182

TABLE OF CONTENTS

NSX-T L2VPN Server Configuration	184
NSX-T L2VPN Session Configuration	185
NSX-T L2VPN Segment Configuration	187
NSX-T L2VPN Supported Clients.....	189
NSX-T L2VPN Peer Compatibility Matrix.....	190
NSX-T Autonomous Edge.....	190
NSX-v Standalone Edge.....	191
NSX-v Managed Edge.....	191
NSX-T L2VPN Client Configuration.....	191
NSX-T L2VPN Session Configuration	192
NSX-T L2VPN Segment Configuration	193
Summary.....	194
Chapter 7: NSX Intelligence, NSX-T Alarms/Events, and Network Visualizations.....	195
NSX Intelligence.....	195
NSX Intelligence Use Cases.....	196
NSX Intelligence vs vRealize Network Insight.....	196
NSX Intelligence Requirements	197
NSX Intelligence Footprint.....	197
NSX Intelligence Deployment	198
NSX Intelligence Verification	202
Flow Visualization with NSX Intelligence.....	202
Recommendations with NSX Intelligence.....	205
Alarms with NSX Intelligence	209
NSX-T Network Topology	210
NSX-T Network Topology Use Cases	210
NSX-T Network Topology GUI.....	210

NSX-T Alarms and Events	212
NSX-T Native Monitoring Tools	212
NSX-T Events.....	213
NSX-T Alarms	213
NSX-T Alarm Use Cases.....	213
NSX-T Alarm and Events Architecture	214
NSX-T Alarm Definitions	215
NSX-T Alarm Definition Configuration	216
NSX-T Alarm Verification	217
NSX-T Alarm Actions	217
NSX-T Alarm Suppression	218
View NSX-T Open Alarms	219
Summary.....	219
Chapter 8: Authentication and Authorization	221
VMware Identity Manager (vIDM).....	221
vIDM with NSX-T Integration	222
VIDM with NSX-T integration Prerequisites.....	222
VIDM with NSX-T Integration.....	224
Creating an OAuth Client	225
SHA-256 Certificate Thumbprint.....	227
VIDM with NSX-T Integration Configuration.....	228
vIDM with NSX-T Integration Verification	231
Logging In with vIDM Services Enabled	231
Default NSX-T GUI Login Screen.....	231
Local Login When VIDM Is Enabled.....	232
NSX-T LDAP Integration	233
LDAP	233
LDAP and NSX-T Integration Benefits.....	234

TABLE OF CONTENTS

- LDAP Authentication and NSX-T 234
- Identity Source Configuration 235
- LDAP Server Configuration 236
- GUI Login Using LDAP 237
- Role-Based Access Control (RBAC) 238
 - NSX-T User Account Types 238
 - Policy Management, Access, and Authentication 238
 - Local NSX-T Users 239
 - Authentication Policy Setting Configuration for VIDM Users..... 241
 - Authentication Policy Setting Configuration for LDAP Users 242
- Summary..... 246
- Chapter 9: NSX-T Federation 247**
 - NSX-T Federation 247
 - NSX-T Federation Use Cases 247
 - Policy Consistency and Simplification of Operations..... 248
 - NSX-T Federation Components 249
 - NSX-T Federation Consumption Methods..... 252
 - LM Configuration Ownership 253
 - GM Configuration Ownership..... 254
 - LM Configuration 254
 - GM (Federation) Configuration..... 255
 - NSX-T Federation Networking 265
 - NSX-T Federation Stretched Networking..... 265
 - Logical Topologies with Tier-0 and Tier-1 Gateways..... 267
 - Tier-0 and Tier-1 Gateway Deployment Rules 269
 - Tier-0 and Tier-1 Single Location Deployments..... 270
 - Tier-0 and Tier-1 Multi-Location Deployments 271
 - Tier-0 Multi-Location Stretched Modes 272

Tier-1 Multi-Location Stretched Modes	274
Remote TEP (RTEP)	276
Stretched Layer-2 Network	277
Stretched Layer-2 Network VNI Mapping	277
Stretched Layer-2 Packetwalk.....	277
NSX-T Federation Security.....	287
NSX-T Federation Security Use Cases.....	287
NSX-T Federation Security Components	288
GM Firewall Policy	288
GM Firewall Rules.....	289
GM and LM Section Overlap	290
Global Groups	291
Regions.....	292
NSX-T Federation Security Tags	293
Federation Security Configuration Steps	294
GM Group Types and Span of Group Types	295
GM Group Rules.....	297
GM Group Span of Dynamic Members.....	298
GM Group Based on a Virtual Machine Tag	299
Summary.....	301
Chapter 10: Public Cloud Integration.....	303
NSX-T on VMware Hyperscalers	303
What Is a VMware Hyperscaler?	303
VMware Hyperscaler Use Cases	306
NSX Cloud.....	309
Summary.....	316

TABLE OF CONTENTS

Chapter 11: Cloud Management Platform Integration and Automation.....	317
VMware Cloud Management Platforms.....	317
vRealize Automation (vRA).....	318
vRealize Orchestrator (vRO).....	319
VMware Cloud Director.....	319
Other Cloud Management Platforms.....	321
Ansible.....	321
Terraform.....	321
PowerShell/PowerCLI.....	322
Summary.....	322
Index.....	323

About the Author



Iwan Hoogendoorn started his IT career in 1999 as a helpdesk agent.

Soon after, Iwan started to learn Microsoft products, which resulted in his MCP, MCSA, MCDBA, and MCSE certification.

While working as a Microsoft Systems Engineer, Iwan gained an excellent basis and developed additional skills and knowledge in computer networking. Networking became his passion in life. This passion resulted in learning networking with Cisco products.

Like most network engineers, one of his "dreams" was to work for Cisco. But before this could happen, he needed to finish his bachelor's degree in ICT, and he completed this in 2009.

Early in 2010, he started working for his dream company, Cisco. After finishing his Master's degree (part-time) in Computer Science at the University of Amsterdam and becoming a CCIE (#13084) in six technology areas, Iwan, as ambitious as he sometimes is, wanted to learn something new—virtualization. Because networking was something that ran through his veins, network virtualization was the next logical step. So he decided to learn VMware NSX.

He got the opportunity to work for VMware in 2016 as a Senior NSX PSO Consultant. During his time at VMware, he gained more knowledge on private and public clouds and the related products that VMware developed to build the Software Defined Data Center (SDDC).

ABOUT THE AUTHOR

After working for four years as a senior NSX PSO consultant (primarily with VMware NSX-v and NSX-T), Iwan was promoted to a staff SDDC consultant, focusing on the full SDDC stack, including hyperscaler offerings on the main public clouds like AWS (VMC on AWS), Microsoft (Azure VMware Solution), and Google (Google Cloud VMware Engine).

Iwan is certified in multiple VMware products, including NSX, and he is actively working with VMware certification to develop network-related exams for VMware. In addition to his VMware certifications, Iwan is also AWS and TOGAF certified.

Iwan is the author of the Apress book *Getting Started with NSX-T: Logical Routing and Switching*.

About the Technical Reviewers



It was an early start for **Abdullah Ibrahim Abdullah**, as he ventured into the workforce at the age of 10. He jumped between computer shops and found his passion and pursued it. Abdullah currently leads business transformation engagements and has been in the industry for more than 15 years, successfully designing and delivering more than 80 projects, locally and internationally, for industry-leading enterprises (with some of them in the Fortune 500). Abdullah currently focuses on VMware’s full stack and delivers

horizontally and vertically across the different VMware technology streams.

Abdullah holds [100+ industry accreditations](#) and is a VMware certified design expert in network virtualization ([VCDX #270](#)). He graduated in his early stages from a vocational institute with a TS (technique superior) in computer systems and networks and has a Bachelor’s degree in business administration (business information systems).

Abdullah is an avid and active community endorser and cherishes sharing knowledge. He has been an [vExpert](#) since 2013, a VMware education SME since 2018, a VMTN Community Warrior (2019), and a member of VMware’s Community Influencer 100 Club.

Abdullah can be reached via Twitter at [@doodzZZ](#) and via [his blog](#).

ABOUT THE TECHNICAL REVIEWERS



Rutger Blom, VCIX-NV/VCIX-DCV, is a senior consultant at Proact IT, focusing on design and implementation of VMware NSX and VMware SDDC. He is a 25-year industry veteran and has worked at Proact IT for five years. He has contracted with VMware PSO on a nearly full-time basis since 2018. Rutger is a VMware certified implementation expert for VMware NSX as well as VMware vSphere and was recognized as a VMware vExpert in 2019, 2020, and 2021. He maintains a blog dedicated

to VMware NSX and has spoken at local VMware User Group meetings. Follow Rutger on Twitter at @rutgerblom or visit his blog at <https://rutgerblom.com>.

Acknowledgments

Thank you VMware, for allowing me to write this book.

Thank you VMware #vExpert community, for the knowledge you've shared.

Thank you Abdullah and Rutger, for working together with me on this book and making it happen.

Thank you Apress, for allowing me to write this book and publish it on my behalf.

Introduction

Chapter 1 explains the challenges found in the traditional data center security model and how NSX-T provides macro- and micro-segmentation to offer a zero-trust security model to overcome these challenges.

Chapter 2 explains what distributed IDS and URL analysis are and covers their use cases. It also explains the architecture of how NSX-T implements distributed IDS and URL analysis.

Chapter 3 explains the differences between the types of service insertion and how the north-south and east-west network sections work. You also learn how endpoint protection works and about its typical use cases.

Chapter 4 explains the NSX-T network services like Network Address Translation (NAT), Dynamic Host Configuration Protocol (DHCP), and Domain Name Services (DNS).

Chapter 5 explains the NSX-T load balancer architecture and discusses its components.

Chapter 6 describes the requirements for IPsec VPN using NSX-T and identifies the different VPN types offered. It also describes how to create and configure Layer-3 IPsec VPN tunnels and Layer-2 VPN (L2VPN) tunnels.

Chapter 7 describes NSX intelligence and its use cases. It explains what the NSX intelligence system requirements are and what you need to deploy the NSX intelligence appliance. It also explains how NSX intelligence performs visualization and recommendation capabilities. It describes NSX intelligence alarms and events so you can identify and prevent failures of your NSX-T environment. It also describes the different use cases for the Network Topology feature.

INTRODUCTION

Chapter 8 describes the purpose of the VMware Identity Manager (also known as VMware Workspace ONE Access) and identifies the benefits of integrating it. You learn how to configure the integration between NSX-T and vIDM and how to verify it.

Chapter 9 explains how to use NSX-T stretched networking and stretched security features across multiple sites via NSX-T Federation.

Chapter 10 explains how NSX-T is used inside the commonly used VMware public cloud offerings (hyperscalers) and the commonly used native public cloud providers.

Chapter 11 explains the VMware products and other products (such as API wrappers) that use the NSX-T API to automate certain tasks. As cloud management platform integration and cloud automation can be very complex topics and complete books can be written about them, I only touch the surface here.

CHAPTER 1

NSX-T Security and Firewalls

This chapter explains the challenges in the traditional data center security model and covers how NSX-T provides macro- and micro-segmentations to offer a zero trust security model that overcomes these traditional data center security challenges.

Traditional Data Center Security

The traditional data center security model is based on environments that are tied to networks. Typical data center network security models offer security between network subnets, but not within network subnets. This means that all systems that are part of the same subnet can openly communicate with each other unless they have some sort of host-based firewall installed (or something equivalent to that, like IP tables). Common services like DNS, DHCP, or NTP can typically be accessed from anywhere in the network without being checked.

Attackers typically target systems that have a low priority and these systems are overlooked when it comes to installing security patches and updates. When these systems are compromised with a traditional data center security model in place, attackers can freely move around between systems and gain access to private data.

Figure 1-1 shows an illustration of the various network and security boundaries a traditional data center network has.

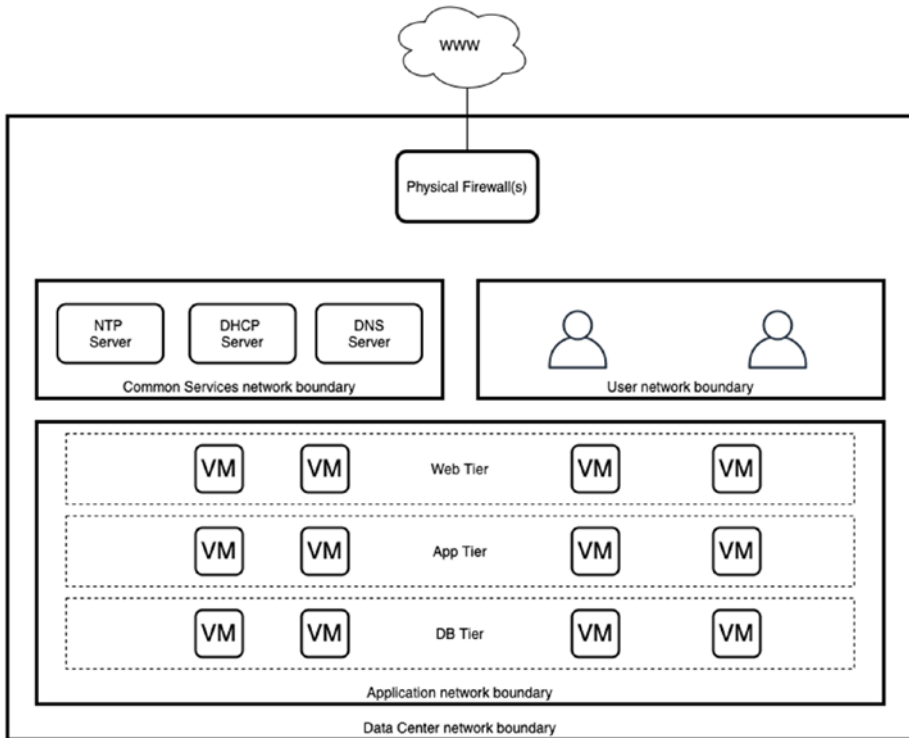


Figure 1-1. Inflexible security boundaries based on subnets

Data Center Security Requirements

In order to overcome the traditional data center security model challenges, the following requirements are introduced:

- Each system (virtual machine or bare metal) will have its own personal firewall and individual security policy.

- In order to flexibly assign a security policy to a system, the security policy needs to be based on several attributes like VM attributes, network attributes, application attributes, or even user attributes, whereby tagging and grouping is used to identify the attributes.
- With applications, additional security controls need to be introduced, such as service insertion, to allow the use of security services/features from third-party vendors and allow agentless virus scanning services.

NSX-T Micro-Segmentation

Micro-segmentation is one of the mechanisms that will satisfy several of the requirements described in the previous section.

With micro-segmentation, you are placing a firewall that is centrally managed by NSX-T in front of each (or selected) virtual machine (or supported physical servers).

With this capability, you are moving from a situation in which your network subnet was your physical boundary and the security policy enforcement was done by a physical firewall northbound, to a situation in which the security policy enforcement is now at the system's (virtual or physical) NIC level.

With this capability, you can logically group your systems in a different and more efficient way and can use a more granular security policy. Each system (virtual or physical) can now have its own security boundary (Figure 1-2)!

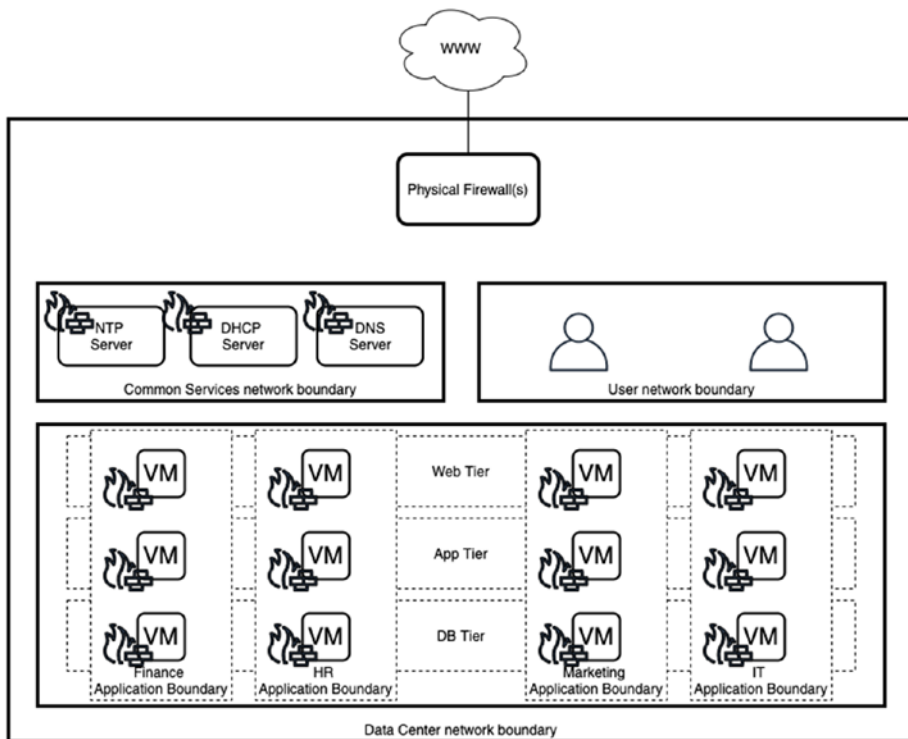


Figure 1-2. Flexible security boundaries based on groups

Zero Trust Security Model

The zero trust security model takes “zero trust” as the main starting point. When we apply this to networks, this means that all systems will block all incoming and outgoing traffic unless that traffic is specifically whitelisted and allowed.

NSX-T Micro-Segmentation Use Cases

You need micro-segmentation to satisfy the following use cases:

- When you need to secure critical applications.
- When you need to secure mobile applications.
- When you need to secure virtual desktop infrastructure.
- When you want to secure systems using agentless endpoint protection solutions.
- When your systems that are part of your DMZ are located across your compute infrastructure (and not necessarily on DMZ islands), whereby logical separation is required rather than physical separation.

NSX-T Micro-Segmentation Benefits

Micro-segmentation with NSX-T is VMware's implementation of the zero trust security model. It does this by leveraging isolation and segmentation and allows you to steer traffic to another virtual firewall using third-party service insertion.

Because zero trust is applied, you minimize the risk of impact whenever there is a network breach.

Micro-segmentation simplifies network flows because you don't have to hairpin the traffic to a northbound firewall. Also, you can use your existing network topology (it's hardware and network topology agnostic).

Another benefit is the mobility of your security policy. When a virtual machine moves from one part of the data center to another part, the security policy moves with it.

Enforcing the Zero Trust Security Model Using Micro-Segmentation

Micro-segmentation can be implemented on multiple levels using different approaches.

- Virtual machine level
- Application level
- Environment level
- Functionality level

I marked the flows in Figure 1-3 with A, B, and C, whereby micro-segmentation takes the following approaches:

- A = Virtual machine level
- B = Application level
- C = Environment level




Figure 1-3. Grouping based on application level

We can also “group” the virtual machines based on their functionality. Figure 1-4 introduces this with B.

- A = Virtual machine level
- B = Functionality level
- C = Environment level

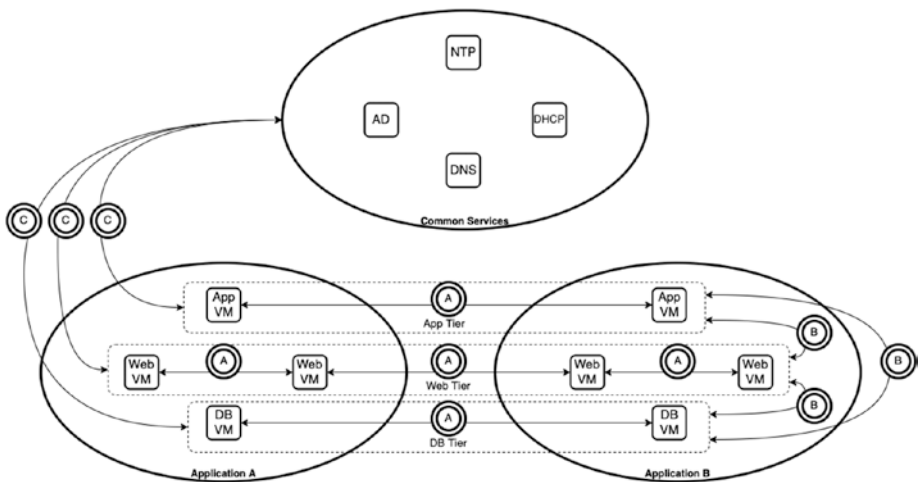


Figure 1-4. Grouping based on functionality level

With these two examples, you can see that you can also logically group systems (virtual machines) and apply security policies to these groups, thereby enforcing micro-segmentation with the use of logically created security groups.

NSX-T Distributed Firewall

The NSX-T distributed firewall (DFW) allows you to implement micro-segmentation to provide zero trust security. The firewall is provided by kernel-based service modules that allow the configuration of firewall rules on the (virtual) NICs of systems networked by NSX-T.

NSX-T Distributed Firewall (DFW) Features

The NSX-T distributed firewall includes the following features:

Availability/recoverability:

- vSphere vMotion support: Firewall policies move with VMs

Manageability:

- Centralized configuration through the NSX UI or REST API
- Support for on-premises and public clouds
- Programmed in the kernel and implemented at the virtual NIC level
- Support for multiple hypervisors
- Support for multiple workloads (VM, bare metal, and container)
- Static and dynamic grouping based on compute objects and tags

Performance:

- Line-rate firewall throughput support

Security:

- Distributed (OSI) Layer 2-4 firewall support
- (OSI) Layer-7 context-aware firewall support
- Firewall rule enforcement regardless of the network transport type (overlay or VLAN)
- Identity-based firewall support (based on active directory users and/or groups)

NSX-T DFW Concepts

When you are dealing with the NSX-T distributed firewall, a few new concepts are introduced that are important for you to understand (Table 1-1).

Table 1-1. *Security Policy Overview*

Concept	Description
Security policy	A collection of firewall rules.
Firewall rule	The instruction that either allows or denies network traffic based on sources, destinations, and network protocols.
Group	A construct that allows you to statically or dynamically group different objects together to form a logical group. This group can then be used in a firewall rule to be the source or destination of the rule.
Service	You specify a combination of a network port and network protocol that can be used inside a firewall rule in order to either allow or deny this specific “service”.
Context profile	Allows you to inspect the Layer-7 application content of network traffic (packets) to allow or deny them.

Security Policy Overview

The NSX-T user interface enables you to configure security policies at multiple levels. To see an overview of all these policies, choose Security ► Security Overview ► Configuration (Figure 1-5). Table 1-2 explains these policies.



Figure 1-5. NSX-T security policy overview

Table 1-2. Security Policy Overview

Policy	Description
Distributed firewall policies	The distributed firewall rule collections that are configured to control east-west network traffic.
Gateway policies	The gateway firewall rules that are configured to control north-south network traffic.
Endpoint policies	The configured guest introspection services and rules.
Network introspection policies	The configured north-south and east-west network traffic redirection rules.
Distributed IDS policies	The configured policies to define intrusion detection rules for east-west network traffic.

DFW Security Policy

The Distributed firewall security policy is a collection of distributed firewall rules that form the overall security policy to control east-west network traffic (Figure 1-6).

NSX-T allows you to group these distributed firewall security policies into several categories, as described in Table 1-3.

Table 1-3. *Security Policy Categories*

Category	Description
Ethernet	This category holds all Layer-2 policies. Layer-2 firewall rules are always evaluated and processed before Layer-3 rules.
Emergency	This category holds all temporary firewall policies that you added in case of an emergency, such as blocking an attacker from attacking a web server. The emergency security policy is evaluated and processed before the Infrastructure, Environment, and Application category rules.
Infrastructure	This category holds all non-application policies specific to infrastructure components, such as vCenter Server, ESXi hosts, and common services like DHCP, DNS, and NTP.
Environment	This category holds all high-level policy groupings. An example of this is that the production group cannot communicate with the testing group or the testing group cannot communicate with the development group.
Application	This category holds all specific and granular application policy rules, such as rules between applications or application tiers, or rules between microservices.

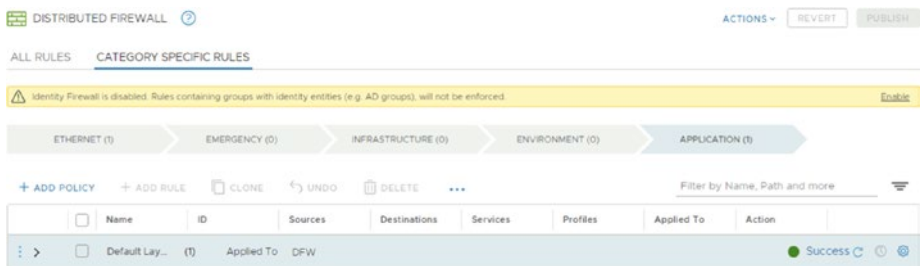


Figure 1-6. *Distributed firewall policy categories*

Each of these categories has its own rules and policies. Firewall rules are enforced left to right and top to bottom.

You can reorder security policies and firewall rules in a specific category. You cannot move policies or rules across different categories.

You can configure your firewall rules under the relevant categories.

DFW Configuration

A firewall policy contains one or more firewall rules in order to either allow or deny specific network traffic. These rules can be stateful or stateless.

A stateful security policy that contains stateful rules allows traffic that has gone out and that comes back in based on the same TCP session. Stateless policies only allow traffic to go out; returning traffic will be denied.

You can create a distributed firewall policy with distributed firewall rules by choosing Security ► East West Security ► Distributed Firewall ► Application ► Add Policy (Figures 1-7 through 1-9).

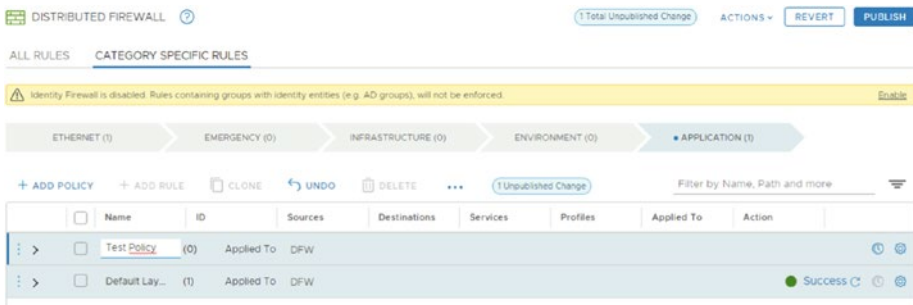


Figure 1-7. Distributed firewall policy configuration (2)

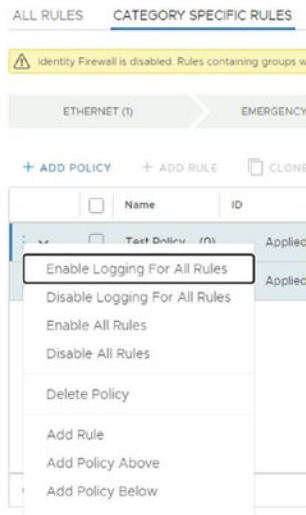


Figure 1-8. Distributed firewall policy configuration (2)




Figure 1-9. Distributed firewall policy rule configuration (1)

In a firewall policy, each firewall rule contains instructions that determine the following factors:

- Whether a packet should be allowed or blocked
- Protocols that the packet is allowed to use
- Ports that the packet is allowed to use

Policies are used for multi-tenancy, such as creating specific rules for finance and HR departments in separate policies.

Firewall rules are enforced in the following ways:

- Like firewall policies, firewall rules are also processed top-to-bottom.
- Each packet is checked against the top rule in the rule table before moving down the next subsequent rules in the table.
- The first rule in the table that matches the traffic’s parameters is enforced immediately. Subsequent rules will not be enforced because the search is terminated for that packet, as it has been enforced by the rule above it.

Because of this behavior, you must place the most granular policies at the top of the rule table.

The default rule at the bottom of the rule table is a catch-all rule. Packets not matching other rules are enforced by the default rule.

After the host transport node preparation operation, the default rule is set to the Allow action. This rule ensures that VM-to-VM communication is not broken during the staging or migration phases.

To implement a zero trust model where only the specified traffic is allowed, the action for the default rule should be changed to block, and firewall policies should be defined to specify all traffic to be allowed (Figure 1-10).




Figure 1-10. Distributed firewall policy rule configuration (2)

DFW Policy Advanced Setting Configuration

When you create a distributed firewall policy, you can specify additional settings, including TCP Strict, Stateful, and Locked (Figures 1-11 and 1-12). Table 1-4 explains these settings in detail.




Figure 1-11. Setting the distributed firewall policy advanced settings



Figure 1-12. Distributed firewall policy advanced settings configuration

Table 1-4. DFW Advanced Settings

Setting	Description
TCP Strict	When you enable this setting, packets that do not complete a three-way TCP handshake are dropped.
Stateful	When you enable this setting, the distributed firewall performs a stateful packet inspection and tracks the state of the network connections. Packets that match a known active connection are allowed (back in). Packets that do not match are inspected against the existing firewall rules.
Locked	When you enable this setting, this allows you to lock a policy while making configuration changes so that other users cannot make simultaneous modifications to the same policy.

DFW Time-Based Security Policy Configuration

When you want your security policy to be active only for a specific time slot, you can configure a time-based security policy. You can do this by clicking the little clock shown in Figure 1-13.




Figure 1-13. Setting the time windows on a distributed firewall policy

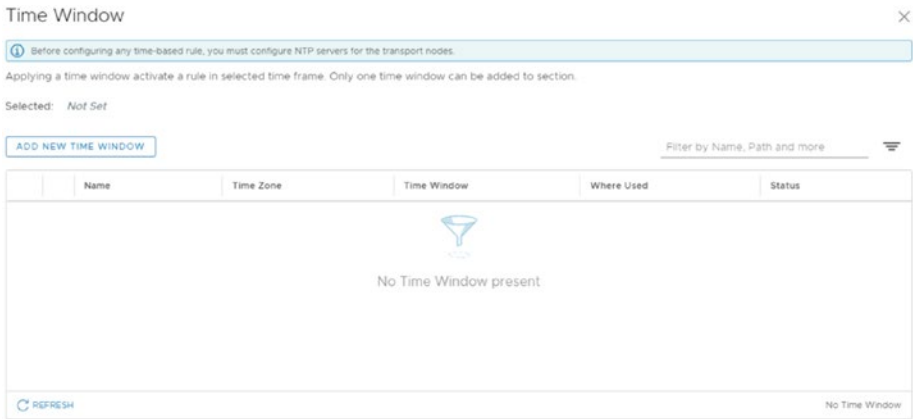


Figure 1-14. Distributed firewall policy time windows summary (without a time window)

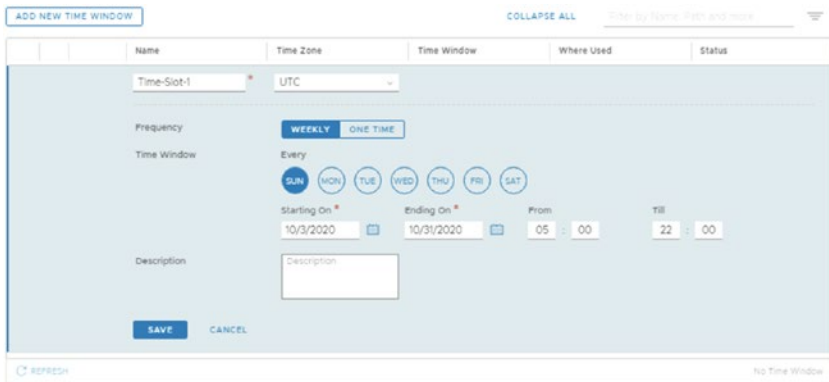


Figure 1-15. Distributed firewall policy time windows configuration




Figure 1-16. Distributed firewall policy time windows summary (with time window)

The parameters to configure a time window are listed in Table 1-5.

Table 1-5. *Time Window Parameters*

Parameter	Description
Name	Time-Slot-1
Time zone	(UTC or local to transport node)
Frequency	(weekly or one time)
Recurring days	Every Sunday
Start and end date	2-OCT-2020 / 31-OCT-2020
Start and end time	05:00 / 22:00

The configured time window activates the security policy between October 2nd and October 31st every Sunday between 05:00 in the morning until 22:00 in the evening.

The From and Till parameters must be configured in 30-minute increments. For example, from 09:00 to 09:30 is a valid configuration. If you configure an interval from 08:15 to 08:45, a configuration error appears in the UI. Before you configure a time-based rule, you must configure NTP servers for the transport nodes.

DFW Rule Configuration

Distributed firewall rules are a set of criteria that are used to evaluate the network traffic flows. They contain the instructions that determine if a packet should be allowed or blocked.

You can configure the rules by clicking the three dots in front of the security policy, as shown in Figure 1-17.

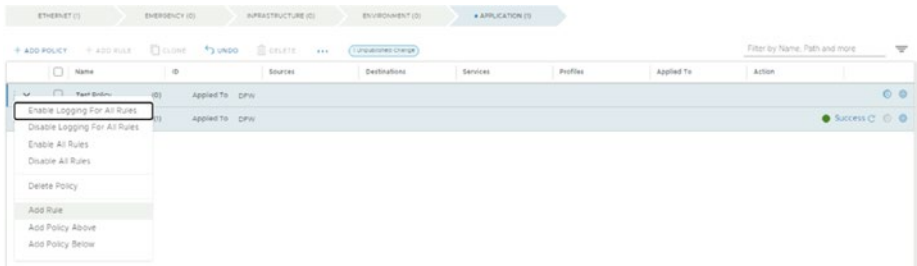


Figure 1-17. Adding a distributed firewall rule

The parameters required to configure the criteria are described in the next section.

DFW Rule Parameters Configuration

To configure a firewall rule, you must include parameters like Source, Destination, Service, Profile, Applied To, and Logging, and you need to select which action should be taken when a rule is matched (or hit). See Figure 1-18.

The screenshot shows a table with columns: Name, ID, Sources, Destinations, Services, Profiles, Applied To, and Action. Three rows are visible:

Name	ID	Sources	Destinations	Services	Profiles	Applied To	Action
Test Policy	(1)	Applied To	DFW				
Test Rule	(1)	Any	Any	Any	None	DFW	Allow <input checked="" type="checkbox"/> <input type="checkbox"/>
Default Layer3 Sec...	(1)	Applied To	DFW				Success <input checked="" type="checkbox"/> <input type="checkbox"/>

Figure 1-18. Distributed firewall rule parameters

The parameters that can be defined when configuring a distributed firewall rule are listed in Table 1-6.

Table 1-6. *Distributed Firewall Rule Parameters*

Rule	Description
Name	You can specify a name for the rule.
Sources	You can use previously defined groups.
Destinations	You can use previously defined groups.
Services	You can specify a port and protocol combination.
Profiles	You use context profiles to define context-aware or Layer-7 rules.
Applied To	Defines the scope of the rule. It can be the full DFW or a specific security group.
Action	You can select from the following firewall rule actions: <ul style="list-style-type: none"> • Allow • Drop • Reject

The order of the distributed firewall rules determines how the traffic is handled. You can drag and drop rules in the UI to change their order.

DFW Source and Destination Definitions

The sources and destinations for a distributed firewall rule can be an IP address (IPv4 or IPv6) or a MAC address. You can also use objects, such as a virtual machine or a group. The default is Any, when you do not make the source or destination specific.

Groups

When you use a group as the source and/or the destination, you are using a collection of assets. Defined assets can be virtual machine objects, virtual interfaces (VIFs), segments, segment ports, IP addresses, IP subnets, MAC addresses, active directory user groups, and physical servers.

Note Before creating a group that includes AD users, you must add an AD domain to NSX Manager.

You add this domain through the NSX-T UI by choosing System ► Configuration ► Identity Firewall AD ► Add Active Directory. The main use case for creating a group that includes AD users is to configure identity-based firewall rules.

Creating a Group and Adding Members

You can create a group through the NSX-T UI by choosing Inventory ► Groups ► Add Group, as shown in Figure 1-19.




Figure 1-19. Adding a new group

Click Set Members to define dynamic (Figure 1-20) or static (Figure 1-21) membership criteria, as explained in Table 1-7.

Table 1-7. Group Membership Criteria

Criteria	Description
Dynamic group inclusion	Dynamic group inclusion for virtual machines can be based on tags, machine names, OS names, or computer names.
Static group inclusion	Static group inclusion criteria apply to virtual machines, virtual interfaces (VIFs), segments, segment ports, IP sets, MAC sets, AD user groups, physical servers, and nested groups.

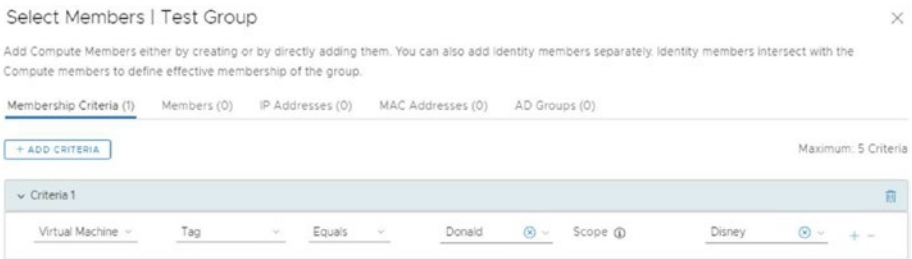


Figure 1-20. Dynamic group membership

Note When you are using nested groups with the static group inclusion, the recommendation is not to use nested groups that are nested with more than three levels.

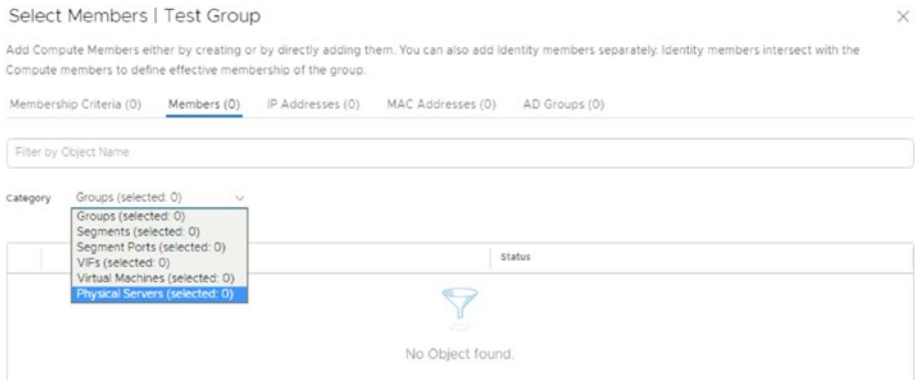


Figure 1-21. Static group membership

Service Specification Inside a DFW Rule

Services are used in distributed firewall rules to identify the types of traffic to block or allow. Services contain the port and protocol definition for network traffic. One or more services can be specified inside a distributed firewall rule.

NSX-T Manager includes a list of predefined services. These services cannot be modified or deleted. You can also create your own services to meet your specific communication requirements.

You can create a service (or view the existing services) through the NSX-T Manager UI by choosing **Inventory** ► **Services** ► **Add Service**, as shown in Figure 1-22.

CHAPTER 1 NSX-T SECURITY AND FIREWALLS

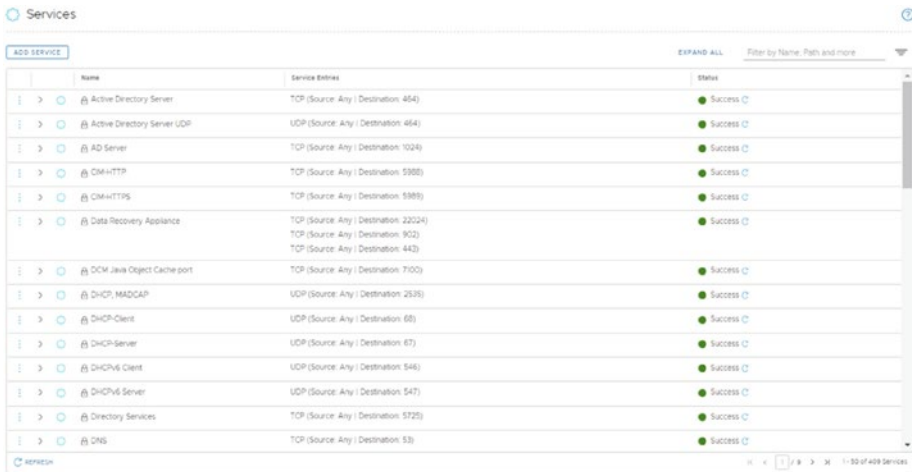


Figure 1-22. Adding a service

You can also add a service (or select an existing service) when you are in the process of configuring the distributed firewall rule. You do this by clicking the pencil in the Services column of the distributed firewall rule itself (Figures 1-23 through 1-25).




Figure 1-23. Adding a service through the distributed firewall rule (1)

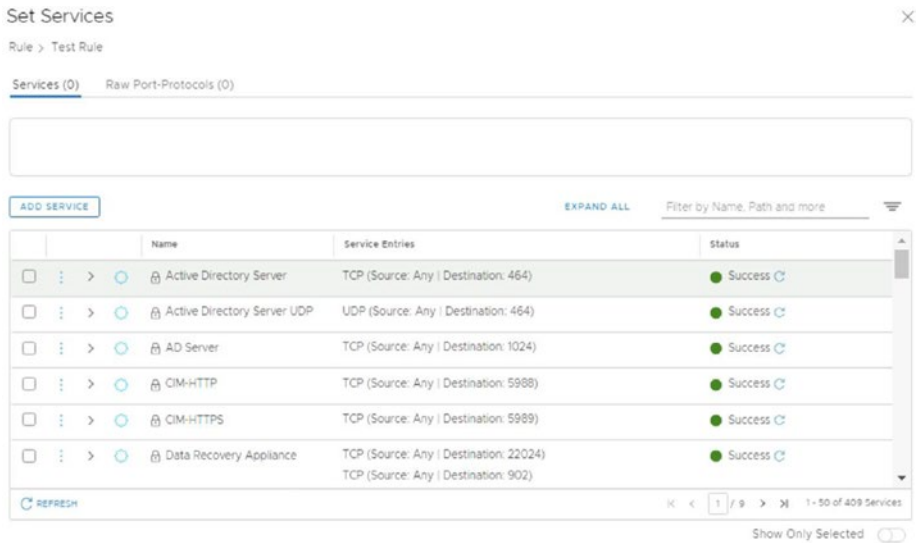


Figure 1-24. Adding a service through the distributed firewall rule (2)




Figure 1-25. Adding a service through the distributed firewall rule (3)

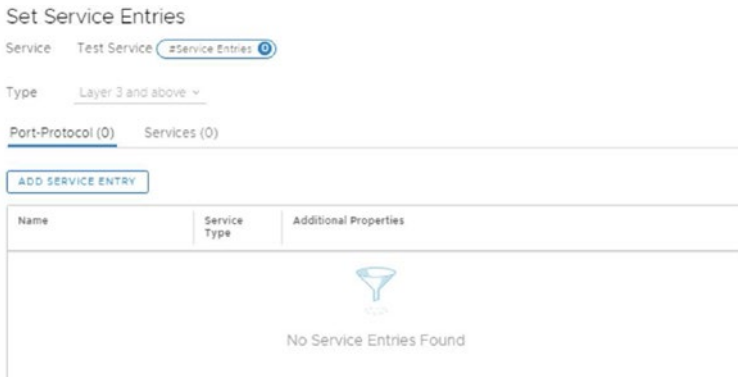


Figure 1-26. Setting the service entries




Figure 1-27. Adding a service entry




Figure 1-28. Selecting the newly added service to distributed firewall rule (1)



Figure 1-29. Selecting the newly added service to distributed firewall rule (2)

Add a Context Profile to a DFW Rule

When you want a distributed firewall rule to inspect packets based on L7 application information, you need to use a context profile.

Just like with services, you can create new context profiles when you create a distributed firewall rule.

NSX-T Manager includes a list of predefined context profiles. You can create your own context profiles to meet your specific communication requirements. Layer-7 firewall rules can be defined only in a stateful firewall policy.

You can create a context profile (or view the existing context profiles) through the NSX-T UI by choosing Inventory ► Context Profile ► Add Context Profile, as shown in Figure 1-30.

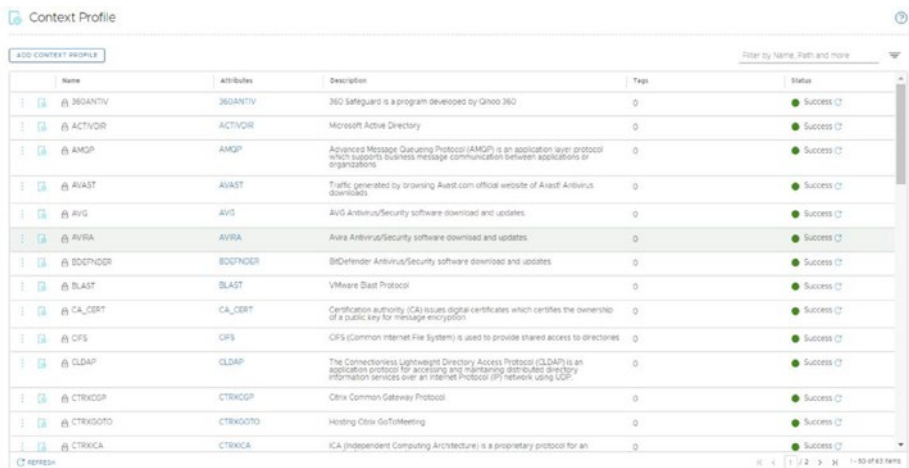


Figure 1-30. Adding a context profile

You can also add a context profile (or select an existing one) when you are in the process of configuring the distributed firewall rule. You do this by clicking the pencil in the profile column of the distributed firewall rule itself (Figures 1-31 through 1-33).




Figure 1-31. Adding a context profile through the distributed firewall rule (1)

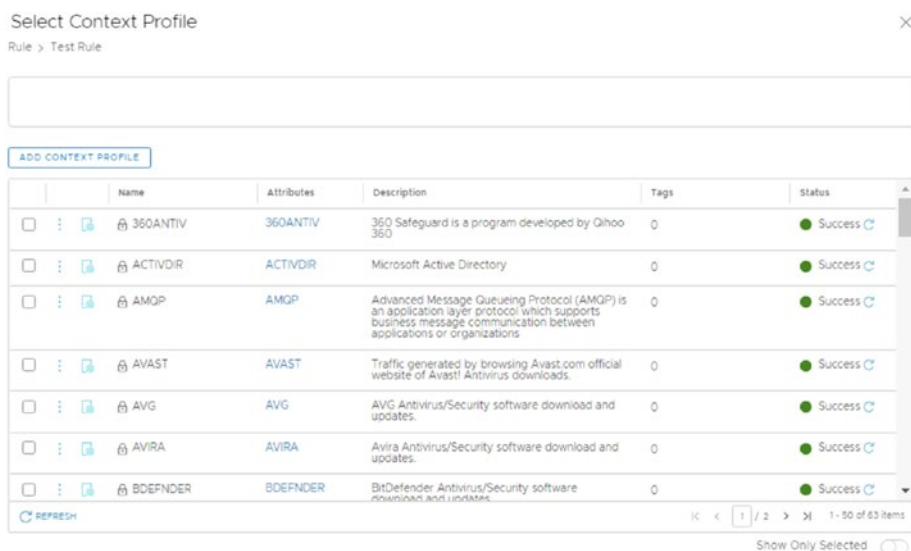
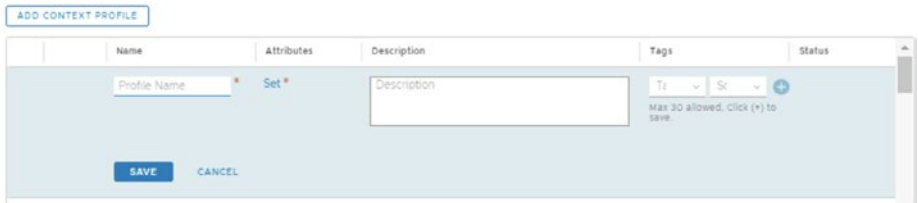


Figure 1-32. Adding a context profile through the distributed firewall rule (2)



Name	Attributes	Description	Tags	Status
Profile Name *	Set *	Description	Tt Sr +	

Max 30 allowed. Click (+) to save.

Figure 1-33. Adding a context profile through the distributed firewall rule (3)

Context Attributes Configuration

A context profile defines context-aware attributes, including application ID and domain name, as well as sub-attributes such as application version or cipher set (Figure 1-34).



Set Attributes

ADD ATTRIBUTE ▾

Search

App Id	Attribute Name/Values	Sub Attribute/Values
<ul style="list-style-type: none"> App Id URL Category Domain(FQDN) Name 		

No Attributes

Figure 1-34. Adding a context profile attribute

Context profiles for distributed firewall rules include the main attributes explained in Table 1-8.

Table 1-8. Context Profile Attributes

Attribute	Description
APP_ID	You can choose from a list of preconfigured applications. You cannot add any applications. Examples include FTP, SSH, and SSL. Certain applications allow users to specify sub-attributes. For example, when choosing SSL administrators, you can specify the TLS_VERSION and the TLS_CIPHER_SUITE. For CIFS, you can specify the SMB_VERSION.
DOMAIN_NAME	You can choose from a static list of fully qualified domain names (FQDNs).

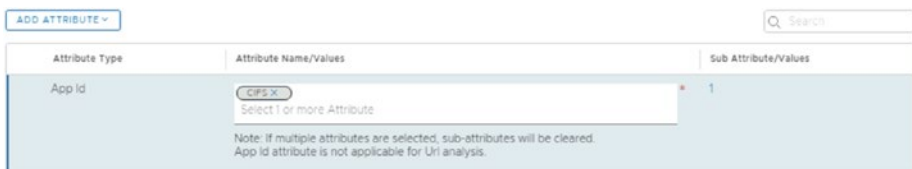
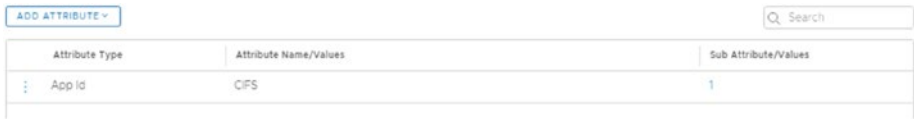


Figure 1-35. Adding a context profile attribute based on app ID (1)

I added a context profile to match the CIFS/SMB v2 traffic, as shown in Figures 1-35 through 1-38.



Figure 1-36. Adding a context profile attribute based on app ID (2)



Attribute Type	Attribute Name/Values	Sub Attribute/Values
App Id	CIFS	1

Figure 1-37. Adding a context profile attribute based on app ID (3)



Name	Attributes	Description	Tags	Status
Test Context Profile	CIFS		0	Success

Figure 1-38. Adding a context profile attribute based on app ID (4)


Now I add an attribute based on a domain (FQDN) name. See Figures 1-39 and 1-40.



Attribute Type	Attribute Name/Values	Sub Attribute/Values
Domain(FQDN) Name	connect.facebook.net	

Note: If multiple attributes are selected, sub-attributes will be cleared. Domain name attribute is not applicable for Uri analysis.

Figure 1-39. Adding a context profile attribute based on a domain (FQDN) name (1)



Attribute Type	Attribute Name/Values	Sub Attribute/Values
Domain(FQDN) Name	connect.facebook.net	

Figure 1-40. Adding a context profile attribute based on a domain (FQDN) name (2)

I added a context profile to match the FQDN connect.facebook.com.

Now that I specified these attributes inside the context profile, I can use them to either allow or block traffic and perform true Layer-7 distributed firewalling.

Scope of the DFW Rule

Using the Applied To attribute in the distributed firewall rule, you can define the scope to which you want the distributed firewall rule applied. You can then ensure that the distributed firewall rule is only applied to the workloads that are in scope of the rule.

When you have three groups—Web, App, and DB—and your firewall rule only contains rules for the Web group as a source and the App group as a destination, it does not make sense to also apply these rules to the DB group (virtual machines), as they will be irrelevant.

This will optimize the resource utilization on the ESXi hosts. It helps to define targeted policies at specific zones or tenants without affecting the policy defined on other zones or tenants.

When you select DFW (the default), the distributed firewall rule will be applied to all the virtual machines in the environment. See Figures 1-41 through 1-43.



Figure 1-41. Defining a scope through the distributed firewall rule



Figure 1-42. Set the scope to DFW

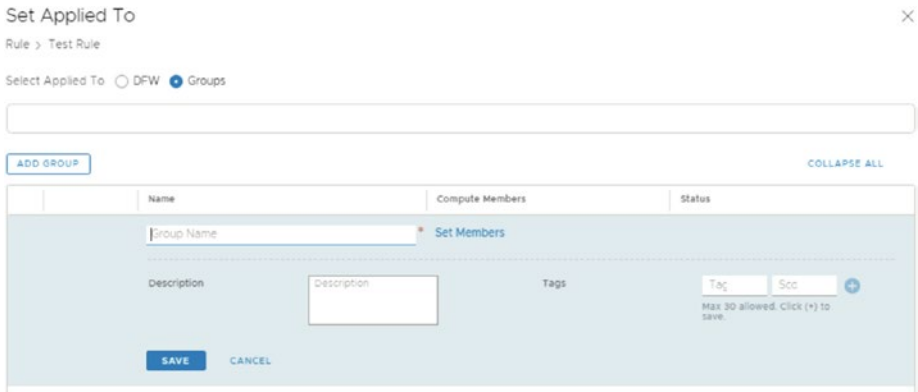


Figure 1-43. Set the scope to specific groups

DFW Rule Actions

You can configure the actions explained in Table 1-9 in a distributed firewall rule. See Figure 1-44.

Table 1-9. DFW Rule Actions

Action	Description
Allow	Allows all L3 and L2 traffic with the specified source, destination, and protocol.
Drop	Drops packets with the specified source, destination, and protocol. Dropping a packet is a silent action with no notification to the source and destination systems.
Reject	Rejects packets with the specified source, destination, and protocol. Rejecting a packet is a more graceful way to deny a packet, because it sends a “destination unreachable” message to the sender.



Figure 1-44. Set a distributed firewall rule action

DFW Rule Settings

You can configure additional settings for your distributed firewall rules, as listed in Table 1-10.

Table 1-10. DFW Rule Settings

Setting	Description
Logging	You can turn logging off or on a per distributed firewall rule basis. The logs are stored in the <code>/var/log/dfwpktlogs.log</code> file on ESXi hosts.
Direction	This setting matches the direction of a packet as it moves across the network. <ul style="list-style-type: none"> • The IN(bound) direction is for traffic ingressing through the distributed firewall. • The OUT(bound) direction is for traffic egressing through the distributed firewall. • The IN_OUT(bound) direction indicates that traffic in both directions is checked.
IP Protocol	The supported protocols are IPv4, IPv6, and IPV4-IPv6.
Log Label	With the use of log labels, you can identify a rule when analyzing the log files.



Figure 1-45. Set the distributed firewall rule settings (1)




Figure 1-46. Set the distributed firewall rule settings (2)

Enable/Disable a DFW Rule

By default, configured distributed firewall rules are enabled. You can disable them with the toggle displayed in Figure 1-47.




Figure 1-47. Enable/disable a distributed firewall rule

Save and View DFW Rule Configurations

You can save and view distributed firewall configurations. When you publish a distributed firewall rule, a draft of the configuration is saved automatically. But you can also save a firewall ruleset at any point in time. See Figures 1-48 through 1-51.

CHAPTER 1 NSX-T SECURITY AND FIREWALLS




Figure 1-48. Publish a firewall rule (1)



Figure 1-49. Publish a firewall rule (2)




Figure 1-50. Save a firewall ruleset (1)

Saved Configurations ×

Fill in the following form and click Save. You can then view the draft under Saved Configurations.

Name

Locked No

Description

Comments

ⓘ A comment is needed to modify the lock on the object

Figure 1-51. Save a firewall ruleset (2)

The Locked switch prevents multiple users from opening and editing a manual firewall ruleset draft. When you enable the lock, you must include a description.

You can then see a timeline with all saved distributed firewall configurations. The types of available saved items are listed in Table 1-11.

Table 1-11. *DFW Rule Saved Items*

Save Item	Description
Auto-saved	Drafts are automatically saved by the system immediately after distributed firewall changes are published by clicking the Publish button. This feature is enabled by default, but can be disabled if you require this.
Saved by others	Distributed firewall configurations are saved by other users different from the user currently logged into the NSX UI.
Saved by me	Distributed firewall configurations saved by the user currently logged into the NSX UI.




Figure 1-52. *Look at a saved firewall ruleset (1)*

Roll Back Saved DFW Rule Configurations

I added a second distributed firewall rule to the security policy, as shown in Figure 1-53.




Figure 1-53. Adding a new distributed firewall rule

If I were to click Publish, the new ruleset would be saved automatically. Say I want to roll back to the previous distributed firewall rule configuration. This requires reverting to the previously published autosave.

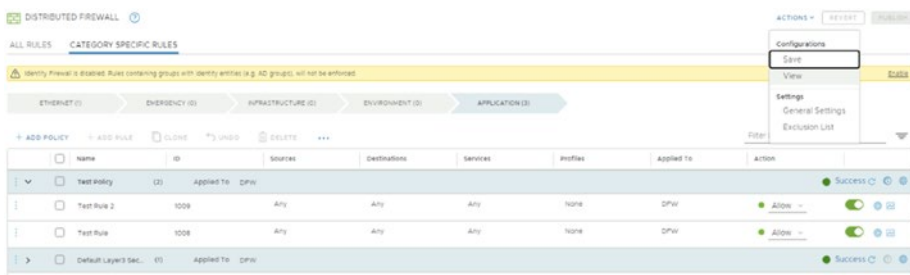


Figure 1-54. View saved distributed firewall rules

In Figure 1-55, you see three bullets. Two of them are dedicated to the autosave, and one of them (the middle) is dedicated to the one I manually saved. You do this by selecting the manual save (Figure 1-56).




Figure 1-55. Look at a saved firewall ruleset




Figure 1-56. Selecting a saved firewall ruleset

This will bring up a screen that displays the details of the distributed firewall configuration on the top, including its name, description, and creation date.

The bottom part of the screen displays the differences between the saved configuration and the last published configuration. See Figure 1-57.

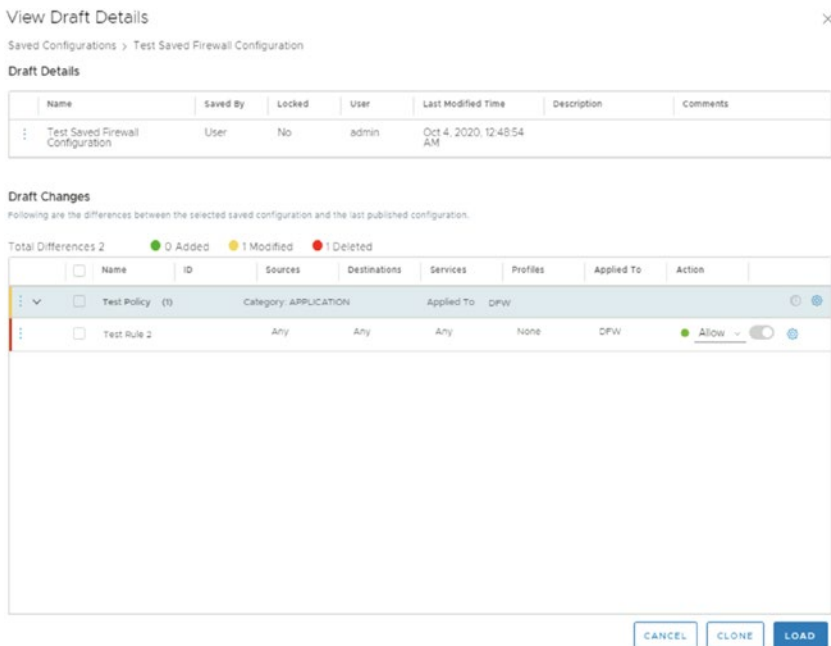


Figure 1-57. View the draft details of the saved distributed firewall ruleset

If you click the firewall ruleset at the point in time you want to revert to, you can then click Load.

Notice the warning:

"You are looking at the saved configuration 'Test Saved Firewall Configuration'. Click Publish to commit or Revert to go back to the last published configuration. You may modify entities as needed."

This means that you first need to publish the ruleset before it is active.

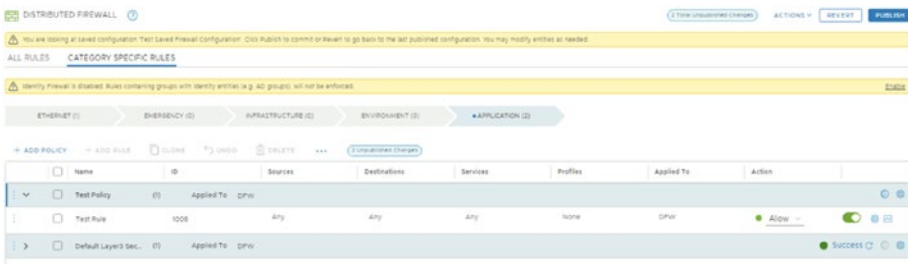


Figure 1-58. Publish the restored distributed firewall ruleset

When you click Publish, you will see a small trash bin in front of the rules that were removed (Figure 1-59).




Figure 1-59. The restored distributed firewall ruleset (after the publish 1)

When you refresh the screen, the rule is gone, as shown in Figure 1-60.

Name	ID	Applied To	Sources	Destinations	Services	Profiles	Applied To	Action
Test Policy	1	Applied To: DFW				None	DFW	In Progress
Test Rule	1008	Applied To: DFW	Any	Any	Any	None	DFW	Success

Figure 1-60. The restored distributed firewall ruleset (after the publish 2)

When you restore a ruleset, you are deleting the existing firewall rules.

This can have an impact on your existing environment. Applications might stop working, or applications might suddenly be accessible that should not be.

Distributed Firewall (DFW) Architecture

A step-by-step high-level overview of the distributed firewall is described in the following steps and shown in Figure 1-61:

1. You configure a distributed firewall policy using the NSX user interface.
2. The configured policy is processed by the NSX-T Manager policy role.
3. The distributed firewall policy is then pushed to the NSX-T Manager role.
4. The NSX-T Manager role forwards the distributed firewall rule configuration to the Central Control Plane (CCP).
5. The CCP then forwards the configuration to the Local Control Plane (LCP) (nsx-proxy) using the Appliance Proxy Hub (APH) with an RPC call using port TCP/1235.

6. The host transport node stores the distributed firewall configuration and configures the data path accordingly.
7. The host transport node sends the rule statistics and status to the NSX Manager with an RPC call using port TCP/1234.

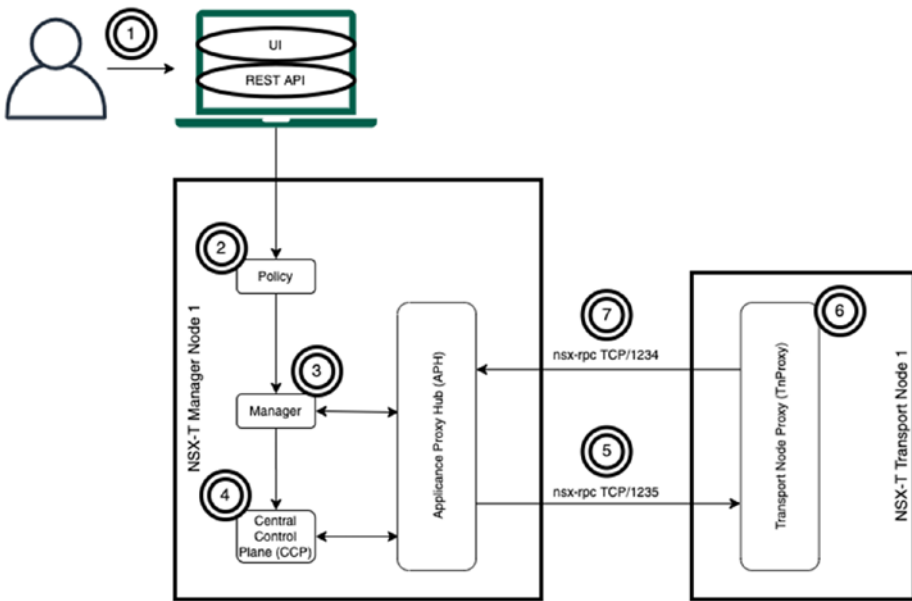


Figure 1-61. The DFW architecture

DFW Architecture on the ESXi Hypervisor

The data path modules outlined in Table 1-12 are responsible for distributed firewall rule processing.

Table 1-12. *ESXi Data Path Modules|DFW Architecture*

Module	Description
VMware Internetworking Service Insertion Platform (VSIP)	The main part of the distributed firewall kernel module that receives the firewall rules and downloads them on each VM's vNIC.
VMware Deep Packet Inspection (VDPI)	A deep packet inspection module daemon in the user space that is responsible for L7 packet inspection. VDPI can identify application IDs and extract context for a traffic flow.

Layer-7 rules (similar to the rest of DFW rules) are programmed into the VSIP module. The VSIP module forwards L7 packets to the VDPI module, which inspects and extracts the Layer-7 information from the packets and returns them to the VSIP module.

The Stats Exporter collects flow records from the VSIP kernel module and generates rule statistics.

A step-by-step high-level overview of how this works follows (see Figure 1-62).

1. The nsx-proxy retrieves the configuration changes from the CCP and configures data path modules.
2. Data path modules:
 - a. The VSIP module receives the firewall rules and downloads them on each virtual machine's vNIC.
 - b. The VDPI module performs the Layer-7 packet inspection.
3. The Stats Exporter collects flow records from the distributed firewall data plane kernel modules and generates rules statistics.

- The nsx-proxy passes the rules statistics and real-time data to the NSX-T management plane with an RPC call using port TCP/1234.

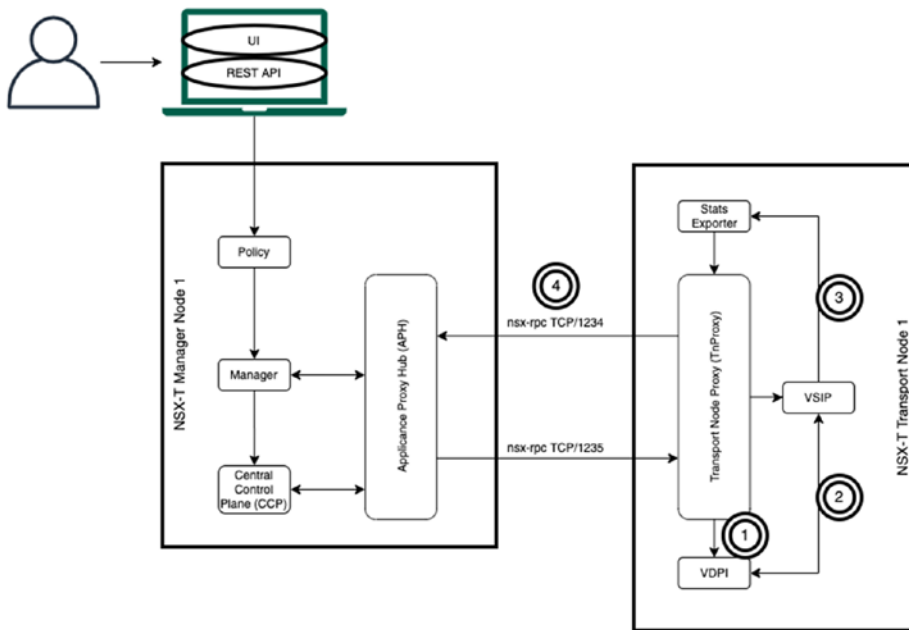


Figure 1-62. The DFW architecture in the ESXi hypervisor

NSX-T Gateway Firewall

The NSX-T Gateway firewall provides stateful (and stateless) north-south firewalling capabilities on the Tier-0 and Tier-1 gateways. This gateway firewall is provided by the NSX-T Edge transport node for both bare-metal and VM form factors. The service router (SR) component provides these gateway firewall services.

The Tier-0 gateway firewall supports stateful firewall filtering only with active-standby high availability mode. The active-active mode supports only stateless firewall rules.

The firewall policy is enforced on the uplink interfaces of the gateways.

This gateway firewall is independent of the distributed firewall rules and can be used together with the distributed firewall.

The Tier-0 gateway service router (SR) component exists on the NSX Edge transport node to enforce the firewall policy before the traffic leaves or enters the NSX-T virtual world. The distributed firewall will perform firewalling for east-west network traffic.

You can configure and perform management on the gateway firewall by using the NSX-T user interface or the REST API framework provided by the NSX-T Manager.

The gateway firewall data path uses the Data Plane Development Kit (DPDK) framework, which is supported on the NSX Edge transport node to provide better throughput.

When you configure the gateway firewall, you will see that it's similar to the configuration of the distributed firewall policy. This configuration is defined as a set of individual rules in a policy. Just like with the distributed firewall, the gateway firewall rules can use logical objects like tagging and groups to build security policies.

Figure 1-63 shows that the Tier-0 gateway and the Tier-1 gateways both have the gateway firewall enabled.

Traffic that is flowing between the 172.16.10.0/24 network toward 172.16.40.0/24 will traverse two Tier-1 gateway firewalls.

Traffic that is flowing between the 172.16.10.0/24 network toward the physical network will traverse one Tier-1 gateway firewall and one Tier-0 gateway firewall.

Traffic that is flowing between the 172.16.10.0/24 network toward 172.16.20.0/24 will traverse only the distributed firewall and no gateway firewall.

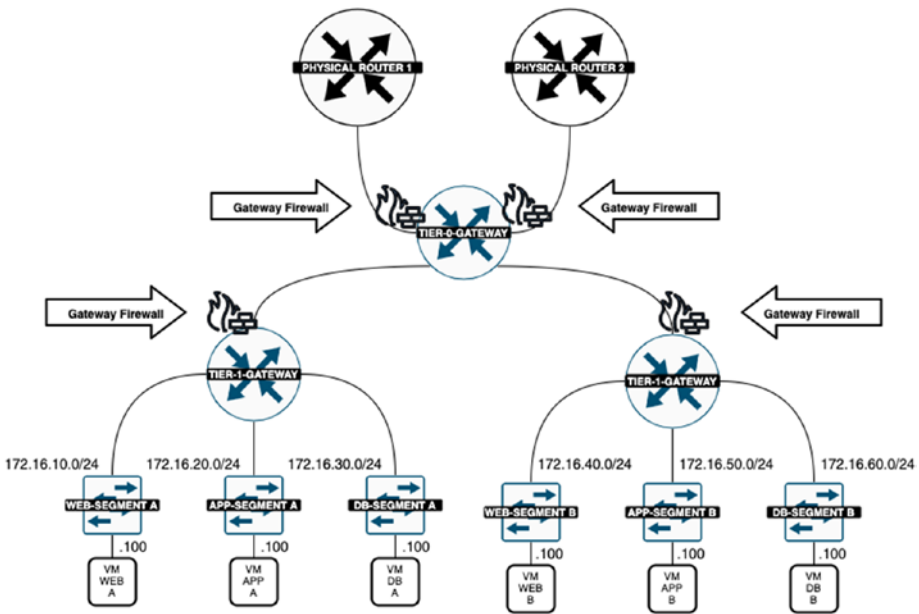


Figure 1-63. Tier-0 and Tier-1 gateway firewalls

Predefined NSX-T Gateway Firewall Categories

Just like with the distributed firewall, the gateway firewall also includes several predefined categories for rules, as described in Table 1-13.

Table 1-13. Gateway Firewall Rule Categories

Rule	Description
Emergency	This category is used for quarantine and can also be used for Allow rules.
System	This category is about hosting automatically generated by the NSX-T data center and specific to internal control plane traffic, such as BFD rules, VPN rules, etc.
Pre Rules	The pre-rules are globally applied across all NSX-T gateways.
Local Gateway	These rules are applied to a particular (specific) NSX-T gateway.
Auto Service Rules	The rules are auto-plumbed rules that are applied to the data plane.
Default	These are the rules that define the default gateway firewall behavior.

The categories are evaluated from left to right.

**Figure 1-64.** Gateway firewall policy categories

NSX-T Gateway Firewall Policy

Just like with a distributed firewall gateway policy, a gateway firewall policy also includes one or more individual firewall rules. These rules are now applied to north-south traffic.

The gateway security policies can be applied to Tier-0 and Tier-1 gateways and their uplink interfaces.

You can configure a gateway firewall through the NSX-T UI by choosing Security ► North South Security ► Gateway Firewall ► Gateway Specific Rules.

Before you start configuring the firewall rules, make sure you select the correct gateway.




Figure 1-65. Gateway specific rules

NSX-T Gateway Firewall Policy Setting Configuration

Just like with the distributed firewall rules (explained in the previous sections), you can also set a time window so that the policy is only applicable during a specified period, and you can configure the advanced firewall policy settings. See Figure 1-66.




Figure 1-66. Gateway policy settings

NSX-T Gateway Firewall Rule Configuration

Configuring a gateway firewall policy and rule is similar to configuring a distributed firewall rule.

You create a policy, then click the three dots and select Add Rule, as shown in Figure 1-67.

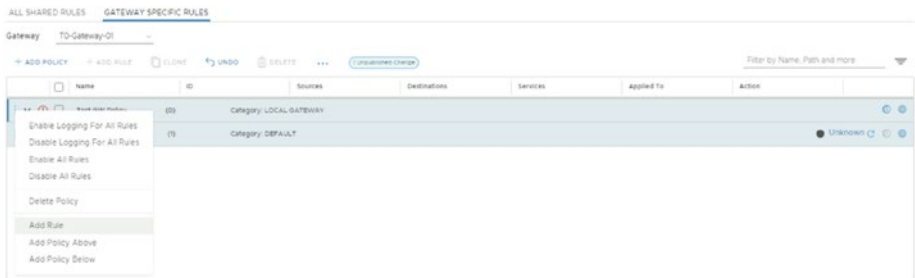


Figure 1-67. Adding a new gateway firewall rule

Once you've created a rule, you can specify a source and destination. This can be one or multiple groups. You also specify a service, where you add a protocol and the desired action that allows, drops, or rejects the packet based on the specified criteria. See Figures 1-68 and 1-69.




Figure 1-68. Specify the new gateway firewall rule parameters



Figure 1-69. Gateway firewall rule action

NSX-T Gateway Firewall Rule Setting Configuration

When you configure a specific rule, you can specify the logging, direction, IP protocol, and log label settings for the gateway firewall rule. Make sure that the firewall rule is published before it takes effect. See Figure 1-70.



Figure 1-70. Firewall rule setting

NSX-T Gateway Firewall Architecture

Let's work through a step-by-step high-level overview of how the gateway firewall works:

1. You configure a gateway firewall policy using the NSX user interface.
2. The gateway policy is processed by the policy role.
3. The gateway policy is then pushed to the NSX-T Manager, which validates the firewall policy and forwards it to the Central Control Plane (CCP).
4. The CCP distributes the firewall configuration through the Appliance Proxy Hub (APH) to the relevant Edge transport nodes.

5. The NSX-proxy receives the firewall configuration from the CCP and configures the edge data path with an RPC call using port TCP/1235.
6. The Stats Exporter collects flow records from the data path and generates the rule statistics.
7. NSX-proxy reports the firewall rules statistics and status to the NSX-T management plane with an RPC call using port TCP/1234 (Figure 1-71).

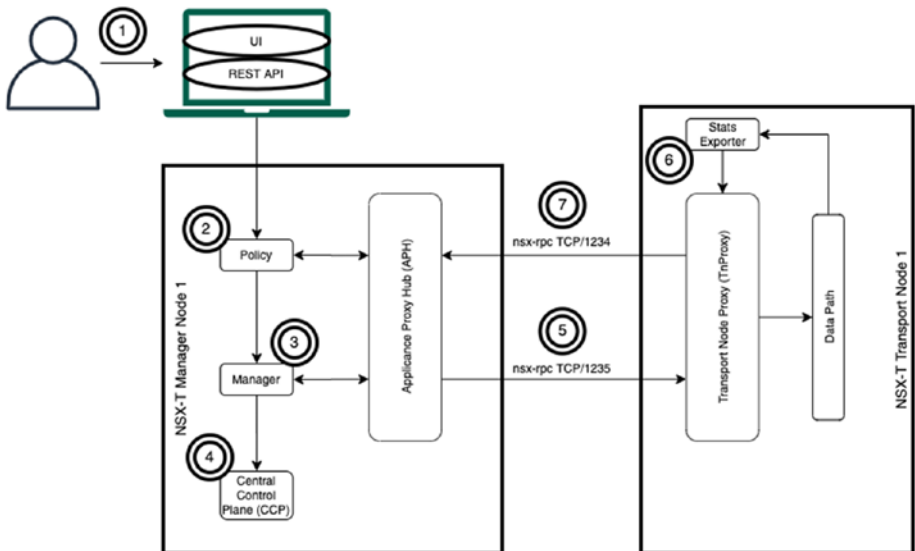


Figure 1-71. NSX-T gateway firewall architecture

Summary

This chapter started with the traditional challenges faced by data center network security, and you learned how NSX-T micro-segmentation can help you with these challenges by using the zero trust security model for configuring the NSX-T distributed firewall (DFW).

You then learned about the capabilities of the gateway firewalls feature that you can configure on Tier-0 and Tier-1 gateways in order to secure north-south incoming and outgoing traffic.

The next chapter explains the advanced security features that NSX-T has to offer, including IDS/IPS and URL analysis.

CHAPTER 2

NSX-T Advanced Security

This chapter explains what distributed IDS and URL analysis are and discusses their use cases. It also explains how NSX-T implements distributed IDS and URL analysis and defines the terminology that you need to know in order to configure them.

Distributed Intrusion Detection

NSX-T uses network introspection to identify malicious intrusion attempts with distributed intrusion detection.

Network introspection is a method that NSX-T uses to send network traffic to another system, either internally or using a third party, to perform an action on the packets. In this case, the action is intrusion detection.

Distributed intrusion detection detects Layer-4 attacks and protects East-West traffic between virtual machines. In order to identify whether traffic is malicious, distributed intrusion detection uses external signatures.

Intrusion detection is distributed across different ESXi nodes (host transport nodes).

Distributed Intrusion Detection Use Case

NSX-T 3.0.x only detects and identifies security vulnerabilities. After the detection, the (security) administrator must manually initiate the quarantine process in order to increase the security level.

Distributed Intrusion Detection Requirements

Before you can use distributed intrusion detection, the distributed intrusion detection components (or VIB¹ modules) need to be installed on each ESXi host (transport node). Distributed intrusion detection should be enabled on standalone hosts and on vSphere clusters. What is also important to know is that the NSX-T Manager nodes need to have access to the Internet in order to download the latest intrusion detection signatures. (When you require a proxy for Internet access settings, this can be provided.)

IDS Signatures

The IDS signatures that are downloaded contain data to identify whether your operating systems that your virtual machines are running have any vulnerabilities. IDS signatures are matched against traffic headers using regular expressions. Currently, NSX-T downloads these signatures from a third-party security repository called Trustwave on a daily basis. Trustwave updates its signatures approximately every two weeks.

IDS signatures are classified in Critical, High, Medium, and Low based on their Common Vulnerability Scoring System (CVSS) score. You can also upload your own signatures.

¹The nsx-idps vSphere Installation Bundle (VIB) provides the intrusion detection system (IDS) features.

IDS Profiles

An IDS profile needs to be configured in order to determine what IDS signatures are included or excluded from detecting vulnerabilities. The default IDS profile is configured to include all signatures marked as critical (Figure 2-1).



Figure 2-1. Default IDS profile

IDS Policy and Rules

Just like with the distributed firewall policy, the distributed IDS policy is a collection of rules, but in this case they are IDS rules.

You need to specify a source, destination, services, IDS profile, and an action to specify where you want to perform intrusion detection. The process, look, and feel works similarly to a distributed firewall (Figure 2-2).




Figure 2-2. Sample IDS profile policy and rule

Distributed Intrusion Detection Architecture

The distributed intrusion detection architecture flow is described like so:

1. The NSX-T Manager downloads the IDS signatures from Trustwave on a daily basis.
2. You configure the IDS profiles and rules based on your requirements.
3. The NSX-T policy component stores the IDS configuration and passes it to the NSX-T Manager component.
4. The NSX-T Manager passes the information to the Central Control Plane (CCP).
5. The CCP pushes the IDS configuration to the ESXi hosts, through the Appliance Proxy Hub (APH).
6. The ESXi hosts store the signature information in NestDB and configure the data path.
7. The ESXi host collects the traffic data and sends events to the NSX-T Manager.
8. The NSX-T Manager parses and displays these events in the UI so you can manually take action if necessary (Figure 2-3).

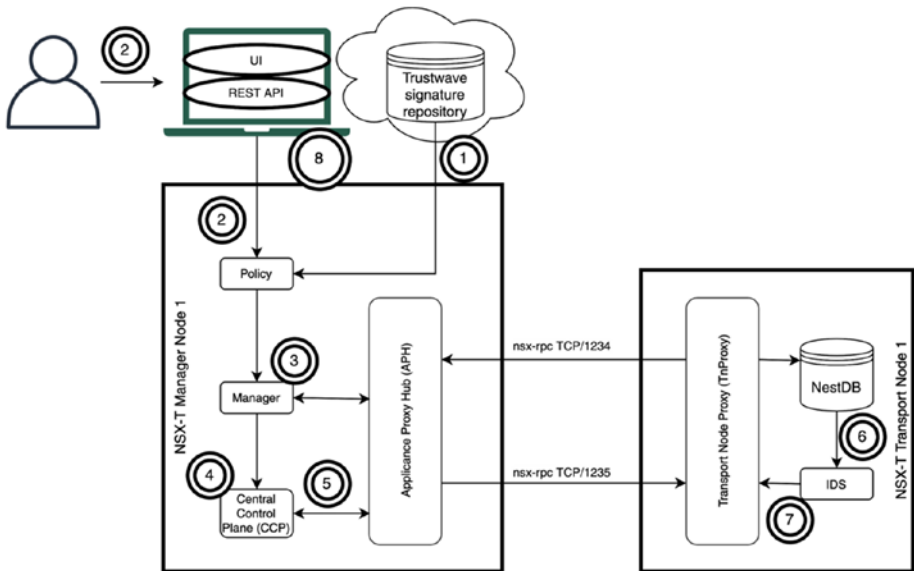


Figure 2-3. Distributed IDS architecture

Distributed Intrusion Detection Configuration

Distributed IDS can be configured on standalone hosts and vSphere clusters. This can be done by navigating to Security ► East West Security ► Distributed IDS ► Settings.

You can also verify when the IDS signatures were last downloaded and determine the version of the signatures (Figures 2-4 and 2-5).




Figure 2-4. IDS signatures before the update




Figure 2-5. IDS signatures after the update

IDS Profile Configuration

A custom IDS profile can be created by navigating to Security ► East West Security ► Distributed IDS ► Profiles.

In this wizard, you configure the IDS signatures that you want to include or exclude for detection based on their severity (Figure 2-6).

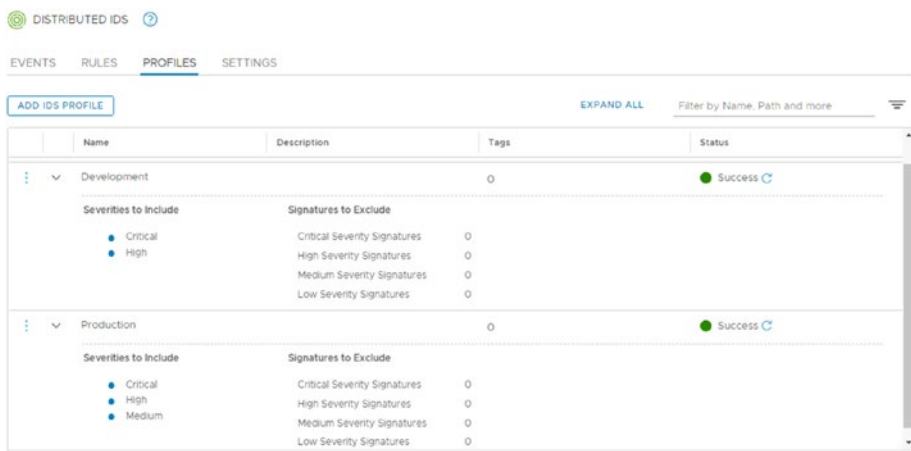


Figure 2-6. IDS profiles

IDS Rules Configuration

IDS rules can be created by navigating to Security ► East West Security ► Distributed IDS ► Rules (Figure 2-7).




Figure 2-7. IDS rules

IDS Event Monitoring

IDS events can be monitored by navigating to Security ► East West Security ► Distributed IDS ► Events (Figure 2-8).

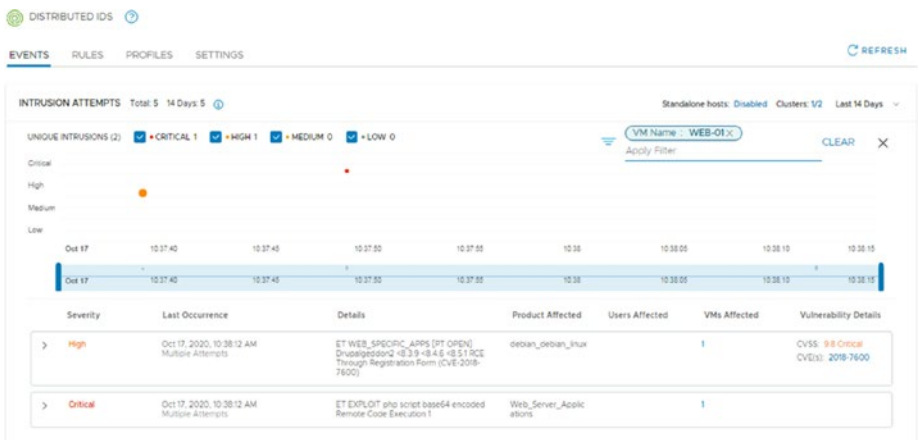


Figure 2-8. IDS events

The IDS Events tab displays all intrusion detection attempts identified in the system. The severity meanings are explained in Table 2-1.

Table 2-1. *Severity Meanings*

Color	Meaning
Red	Critical severity signature events
Orange	High severity signature events
Yellow	Medium severity signature events
Gray	Low severity signature events

You can filter events based on their severity and you can also use freeform text to further filter events.

IDS events are graphically represented using a histogram. You can specify the period that you are interested in by adjusting the blue vertical lines.

Each IDS event type is represented by a dot in the histogram. The size of the dot is proportional to the number of occurrences of that event.

Additional information about each type of event is displayed in a tabular format.

Each occurrence can be expanded to retrieve details about the intrusion attempt, including the attacker, victim, protocol, attack type, and so on.

In this wiki article, I explain how to configure distributed IDS in a more detailed way: https://nsx.ninja/index.php/NSX-T_IDS_end-to-end_testing.

URL Analysis

URL analysis analyzes the North-South traffic, whereby the traffic to external websites is monitored. Websites on the Internet are classified into categories and get a reputation score based on their domain names.

URL Analysis Use Case

URL analysis can be used to get insight into which URLs are accessed in your organization; you can learn more about the reputation of these URLs as well. Based on this information, you can take measures.

URL Analysis Requirements

Before you can use URL analysis, you must perform the following actions:

- The feature must be enabled at the NSX-T Edge cluster level.
- Your edge transport nodes must have access to the Internet to download the category and reputation data that you validate your URLs against.

In order to capture DNS traffic, you also need to enable the Layer-7 gateway firewall on the Tier-1 gateway's uplink interface.

URL Categories

A URL category defines the classification of the URLs. Currently, there are more than 80 predefined categories. A site or domain can belong to multiple categories. For example, www.vmware.com belongs to the Business and Economy category and the Computer and Internet Info category. A domain can belong to a maximum of five categories.

Reputation Score and Severity

The reputation score is a measure of the trustworthiness of a URL; it ranges from 1 to 100.

Based on their reputation scores, URLs are classified into the severities, as listed in Table 2-2.

Table 2-2. *Severity Scoring*

Severity	Numbers	Description
High Risk	1 through 20	Sites with a high probability of containing malicious links or payloads.
Suspicious	21 through 40	Sites with a higher than average probability of containing malicious links or payloads.
Moderate Risk	41 through 60	Benign sites that exhibit some characteristics that suggest security risks.
Low Risk	61 through 80	Benign sites that rarely exhibit characteristics that expose the user to security risks.
Trustworthy	81 through 100	Well-known sites with strong security practices.
Unknown		An unknown reputation score is reported in the NSX UI when the NSX Edge node cannot retrieve the reputation score for a given domain in a timely manner.

URL Analysis Use Case

URL analysis can provide insight into which URLs are accessed inside your organization.

NSX-T will assign a risk score and category so that you can easily identify if the websites accessed by your organization form a threat. Based on the visibility provided by NSX-T, you can also use URL analysis to take proper actions if required.

Webroot as the Cloud Service

NSX-T URL analysis uses Webroot as its source to retrieve metadata on websites.

Webroot is a U.S.-based company that classifies and scores over 95% of the Internet to generate an URL database. Webroot uses a combination of global threat sensors, machine learning algorithms, and human classification. To keep this database updated, Webroot updates the latest knowledge of URL classifications and risk scores for integration on a continuous basis (Figure 2-9).

VMware maintains a cloud service that obtains the reputation and category data from Webroot and caches it for use in the NSX environment.




Figure 2-9. *Webroot and cloud service*

URL Analysis Architecture

URL analysis works as follows:

1. The cloud service retrieves the reputation and category data from Webroot and caches it.

2. The NSX-T Manager obtains the SSL keys from the cloud service and sends the keys to the NSX Edge transport nodes through the Central Control Plane (CCP). Any other configuration changes are also pushed to the NSX Edge transport nodes through the CCP.
3. The NSX Edge transport nodes connect to the (VMware-managed) cloud service using an SSL connection to download the URL category and reputation data locally (every five minutes).
4. The NSX Edge transport nodes then collect all the traffic data for the URLs and send it to NSX-T Manager (every five minutes). The NSX-T Manager then graphically displays the information (Figure 2-10).

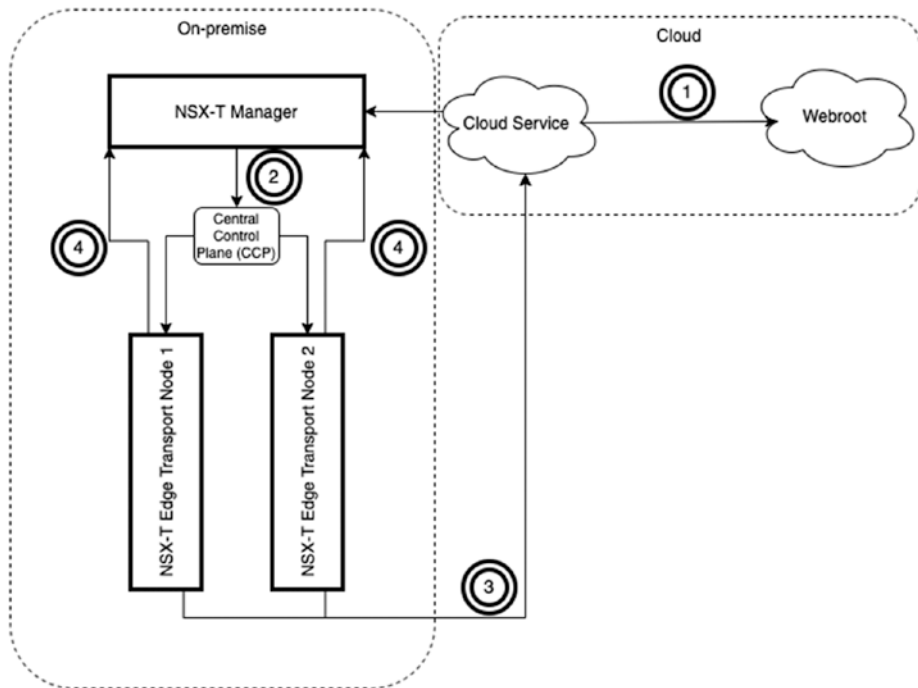


Figure 2-10. URL analysis architecture

URL Analysis Configuration

To configure URL analysis, you need to enable it, create a context profile, and create a Layer-7 gateway firewall rule.

To enable it, you need to navigate to Security ► URL Analysis ► Settings.

You enable URL analysis at the NSX Edge cluster level. When enabled, you can check the connection status to the cloud service for each NSX Edge node. You can also verify the version of the URL data retrieved from the cloud service/Webroot (Figure 2-11).

Edge Cluster / Node Name	Profiles	URL Data Version	Last Synced	Connection Status	URL Analysis State
Pod100-10-Edge-Cluster-01	Set				Enabled

Figure 2-11. Enable URL analysis

URL Analysis Context Profiles

After you enable URL analysis, you can click Set in the Profiles column to configure/create a new context profile.

You can create a context profile in which you select specific URL categories that you want to analyze. In the following example, I selected all available categories. An NSX Edge cluster can have multiple profiles attached.

To set a context profile for URL analysis, follow these steps:

1. Click the Set link on the URL Analysis Settings tab.
2. Add a new context profile and enter a name.
3. Configure the attributes for the context profile.

- The only attribute type available during the URL analysis configuration is URL Category (Figure 2-12).

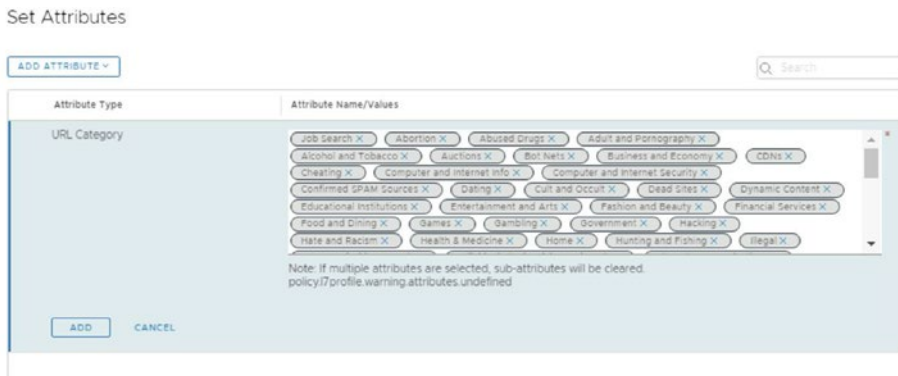


Figure 2-12. Add URL categories to an attribute and context profile

Layer-7 Rule for DNS Traffic

The URL analysis feature relies on the configuration of the Layer-7 gateway firewall rule to capture the DNS traffic that passes the NSX-T Edge cluster.

This Layer-7 gateway firewall rule must be configured on a Tier-1 gateway that is backed by the NSX Edge transport node cluster for which you want to analyze traffic.

The DNS traffic is analyzed and will extract the hostname and IP information from the DNS packets. This information is then used to categorize and score the traffic based on the hostname and IP information.

The services in the gateway Layer-7 rule are set to DNS-UDP and DNS to capture DNS traffic over both TCP and UDP. In addition, the context profile for this rule is set to DNS (Figure 2-13).




Figure 2-13. Tier-1 gateway DNS firewall rule

URL Analysis Dashboard

The URL analysis dashboard (Figure 2-14) shows a summary of all the analyzed URLs classified by reputation score and category. URL categories are also displayed in different colors (Table 2-3).

Table 2-3. Tier-1 Gateway DNS Firewall Rules

Color	URL Category
Red	High Risk URLs
Orange	Suspicious URLs
Yellow	Moderate URLs
Gray	Low Risk URLs
Green	Trustworthy URLs



Figure 2-14. URL analysis dashboard

The left side of the dashboard shows a diagram of the URL distribution. At the bottom of the page are additional details about each URL, including the reputation score, URL, URL category, and the session count.

In this wiki article, I explain how to configure distributed IDS in a more detailed way: https://nsx.ninja/index.php/Configure_NSX-T_URL_Analysis.

Summary

This chapter explained the architecture of distributed IDS and URL analysis. You learned about the use cases, the terminology, and how to configure distributed IDS and URL analysis.

The next chapter explains the differences between the types of service that insertion NSX-T supports in relation to North-South and East-West network traffic.

CHAPTER 3

NSX-T Service Insertion

When you need additional security features that NSX-T does not support out-of-the-box, you can use NSX-T service insertion to redirect network traffic using network introspection to a third-party vendor, or you can use agentless antivirus scanning via endpoint protection.

This chapter covers the differences between the types of service insertion and explains how the north-south and east-west network sections work. You also learn how endpoint protection works and what the typical use cases are for it.

Service Insertion

NSX-T supports network introspection and endpoint protection.

Endpoint protection examines the inside of the guest virtual machines, whereas network introspection examines the outside of the guest virtual machines.

Endpoint protection and network introspection provide internal (endpoint) and external (network) inspection on the activities that are performed by virtual machines. Information about these activities is used by features/services provided by VMware or by third parties like SpoofGuard, identity-aware functions, antivirus solutions, or intrusion detection and protection systems.

Partner services typically provide advanced security features, such as IDS, IPS, L4-L7 firewall, and URL filtering.

Endpoint Protection Use Cases

Endpoint protection handles security on the virtual workload. It enables use cases like agentless antivirus/antimalware, where the agent does not need to be installed on each workload. NSX-T can intercept file events and pass them to a partner virtual appliance that will perform the virus scanning. With endpoint protection, you do not need another agent, and this prevents the processing overhead associated with scanning operations at every workload.

Network Introspection

Network introspection handles data that travels across the network. You can configure redirection rules, which specify which traffic should be inspected by the partner services with features like IDS, IPS, and next-generation firewall.

NSX-T supports north-south and east-west service insertions for network introspection.

Network introspection enables third-party services to examine north-south or east-west traffic passing through a gateway and to take appropriate action.

Service insertion for network introspection can be applied at Tier-0 and Tier-1 gateways to check north-south and east-west traffic.

Using the provided framework and APIs, third-party technology partners can integrate their network and security features/solutions into NSX-T.

North-South Network Introspection

The north-south partner security service is inserted at the data center perimeter between tenant boundaries or between containers.

The insertion points are the uplinks of the Tier-0 or Tier-1 gateway.

The L2 north-south insertion mode (also known as the bump in the wire mode) of the partner service is supported. A partner service virtual machine (SVM) is deployed close to the NSX Edge node to process the redirected traffic.

A service virtual machine (SVM) is a virtual appliance provided by the partner service and is connected over the service plane to receive redirected traffic. Policy-based routing redirection is used for north-south traffic. See Figure 3-1.

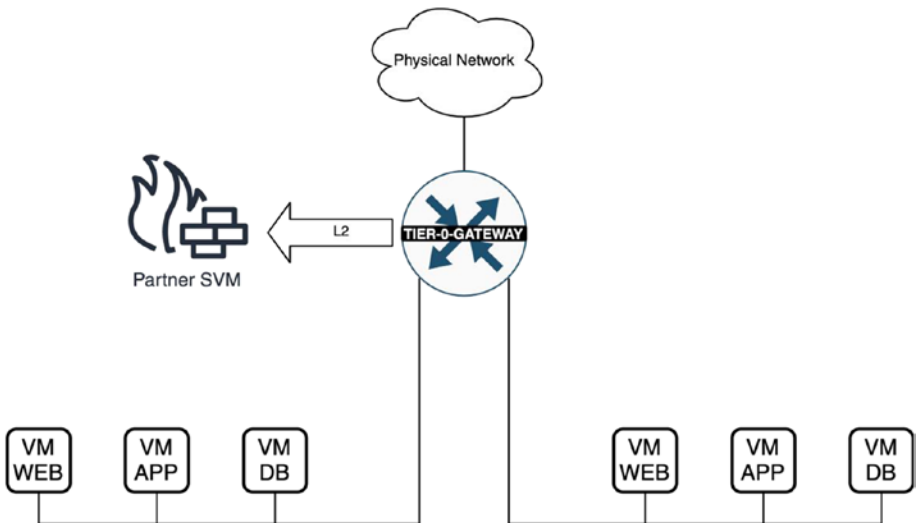


Figure 3-1. North-south partner security

East-West Network Introspection

Partner service virtual machines (SVMs) for east-west network introspection can be deployed either on host transport nodes or in a service cluster.

Common use cases are for VM-integrated next-generation firewalls and micro-segmentation.

Insertion points are at each guest virtual machine’s virtual network card (vNIC).

For SVMs that are deployed on the host transport node, an SVM does not need to be installed on every host.

You may prefer to deploy the partner SVM on each host to achieve the least amount of traffic hair-pinning.

When the partner SVM is deployed in a service cluster, traffic is sent from the compute hosts across the overlay to the hosts in the service cluster.

With east-west service insertion, the insertion points are at each guest virtual machine’s virtual network card (vNIC), so traffic is intercepted at the vNIC of each guest VM. See Figures 3-2 and 3-3.

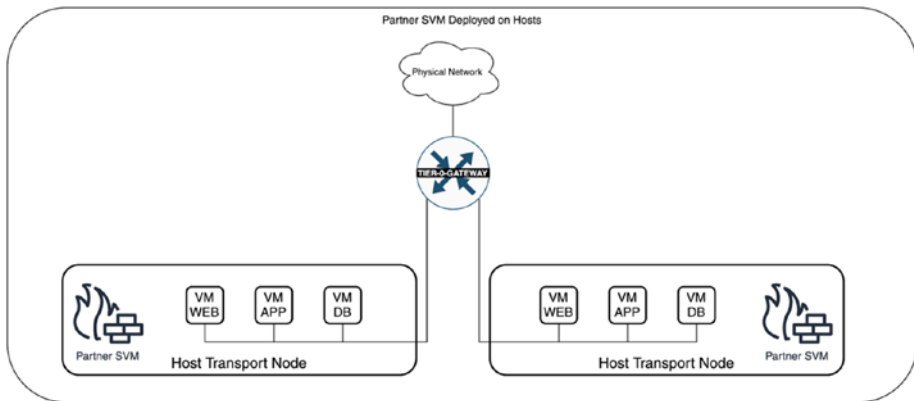


Figure 3-2. East-west partner security on hosts

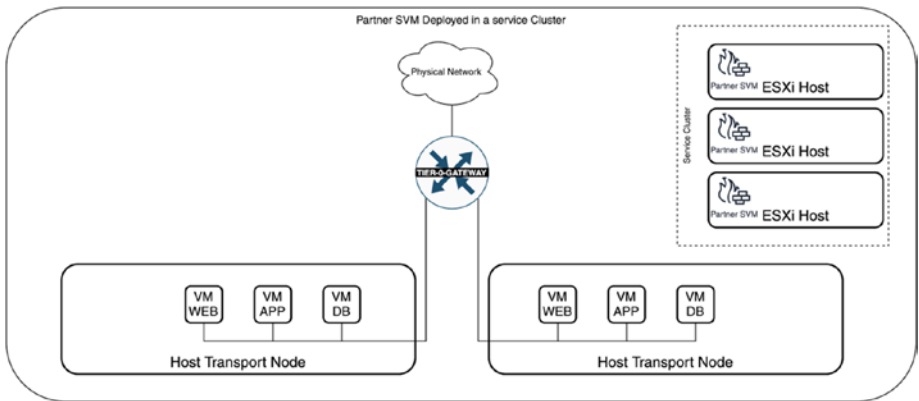


Figure 3-3. East-west partner security on service cluster

Network Introspection Configuration

The configuration steps for north-south or east-west network introspection are similar, as outlined in Table 3-1.

Table 3-1. Network Introspection Configuration Steps

Step	Description	Detailed Description
1	Service registration	This makes the service available to the catalog. <ul style="list-style-type: none"> The partner uses the NSX-T API to register its services to NSXT.
2	Service deployment	This deploys an instance of the registered service. <ul style="list-style-type: none"> East-west: Specify the number of service VMs you want to deploy and the service plane they will connect to. North-south: Select the gateway and host to deploy the service VMs on. Deploy the service VMs using an OVA and vCenter server.

(continued)

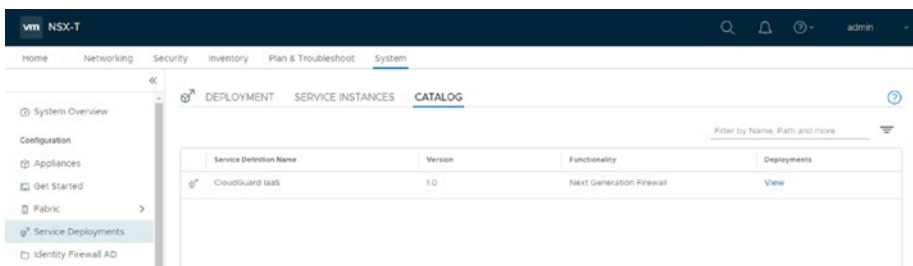
Table 3-1. (continued)

Step	Description	Detailed Description
2	Service consumption	<p>Use services specified in the service chain and redirect traffic.</p> <ul style="list-style-type: none"> • Create a service profile. • Create a service chain that uses service profiles for different services. • Create a redirection policy with steering rules.

Service Registration

A service registration is made by a partner when it makes an API call using the partner management console or CLI. This registration can be automated. To register, you need some parameters including the location of the OVF (URL).

When the service registration is successful, you can see this by navigating to the catalog of the service deployments. Choose System ► Configuration ► Service Deployments ► Catalog, as shown in Figure 3-4.

**Figure 3-4.** Service registration catalog

Service Deployment

After a service is registered and appears in the catalog, you deploy an instance of the service so that it can start processing network traffic. Each partner releases its own partner service OVF for NSX-T integration.

When you deploy a partner service, you need to specify values for several settings, including the service deployment name, attachment points, compute manager, cluster, datastore, and networks.

The deployment process might take some time, depending on the vendor's implementation. The deployment and operational status is monitored so you can keep track of this in the NSX UI. When this is finished and successful, the status will appear as Success.

For north-south service deployments, the following factors must be considered:

- The service instance must be deployed on an ESXi host transport node because logical switching is needed to bridge traffic to the interface where the partner service is attached. This ESXi host must also be managed by vCenter server (i.e., the registered compute manager).
- The partner service instance is deployed on the same host as the edge node to avoid hairpinning situations.
- Each SVM can be applied to only one gateway.

For east-west service deployments, the following factors must be considered:

- From the Service Segments drop-down menu, select or create a service segment. A service segment is associated with a transport zone and is used for the service plane. Guest VMs connected to the service segment are provided east-west network traffic protection.

- From the Deployment Type drop-down menu, select one of the following deployment options:
 - **Clustered:** Deploys the service on a host or hosts belonging to a cluster that is dedicated to host service VMs. You need to specify the number of SVMs required.
 - **Host-Based:** Deploys the service on all the hosts within a cluster.

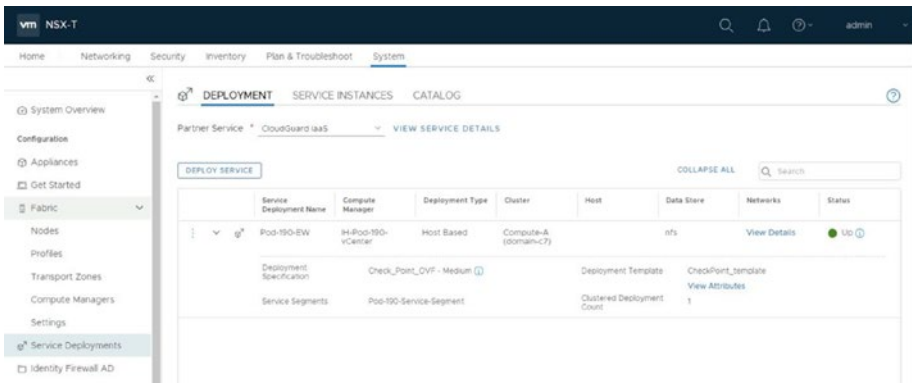


Figure 3-5. Service deployment

Service Consumption

The service consumption consists of the following steps:

- Create a service profile.
- Create a service chain that uses service profiles for different services.
- Create a redirection policy with steering rules.

Service Profile Creation

This is a specific instantiation of a vendor template. If a vendor template defines an IPSec tunneling operation, a service profile specifies details about the IPSec tunnel endpoints, algorithms, etc.

Administrators or third-party vendors can create service profiles. Choose Security ► Settings ► Network Introspection Settings ► Service Profiles ► Add Service Profile, as shown in Figure 3-6.

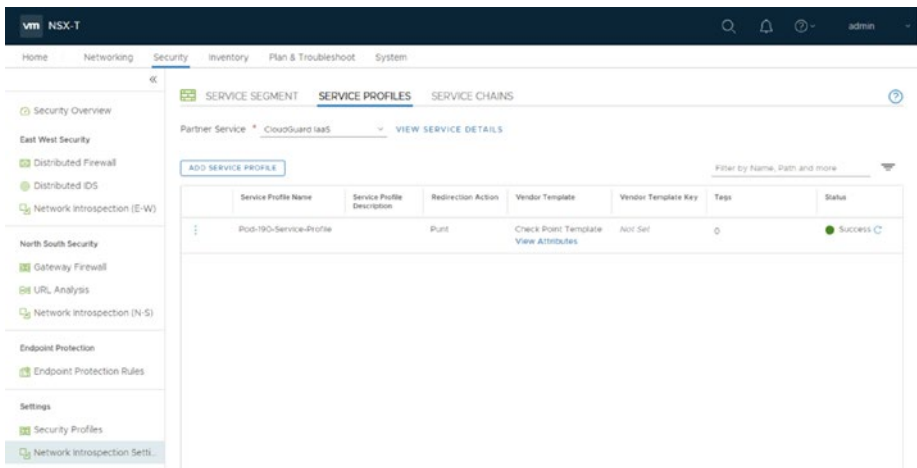


Figure 3-6. Service profiles

Service Chain Creation

This is a sequence of service profiles defined by you. A service chain defines the logical sequence of operations to be applied to network traffic.

One method is to process the firewall-related actions first, then monitor, and then IPSec VPN.

Service chains can specify different sequences of service profiles for egress or ingress traffic directions. Choose Security ► Settings ► Network Introspection Settings ► Service Chains ► Add Chain, as shown in Figure 3-7.

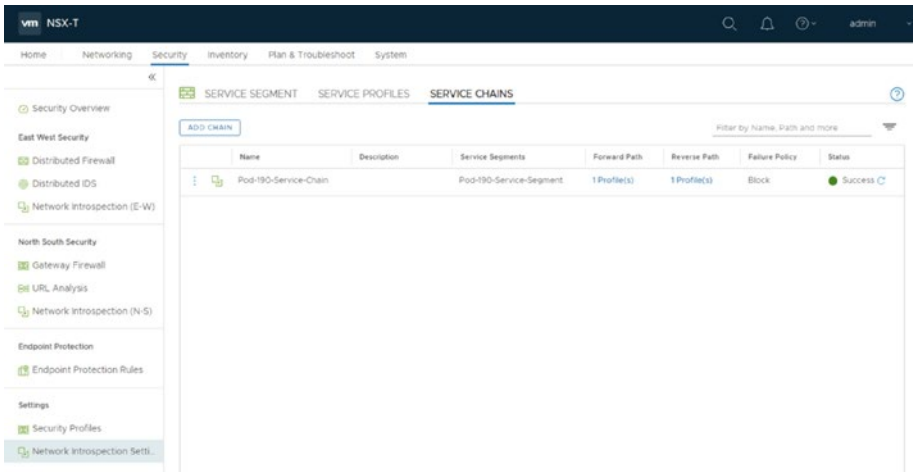


Figure 3-7. Service chain

Redirection Policy Configuration

This is a construct that specifies the traffic matching patterns (based on NSX-T security groups) and a service chain.

All traffic matching the pattern is redirected to the service chain. The redirection policies can be created for north-south and east-west traffic flows.

Figure 3-8 shows an example of the east-west redirection policy. Choose Security ➤ East West Security ➤ Network Introspection (E-W) ➤ Add Policy.

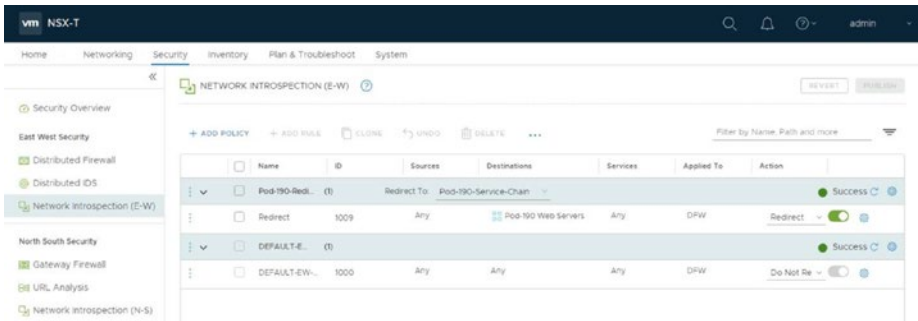


Figure 3-8. Redirection policy

Endpoint Protection Overview

Endpoint protection gives you visibility into guest virtual machine events by analyzing the operating system and file data.

Endpoint Protection Use Cases

The guest introspection capabilities that NSX-T supports include the integration and offloading to third-party services of file data and policy-based agentless antivirus and antimalware protection for Windows and Linux workloads running on vSphere.

Endpoint Protection Process

The endpoint protection process has the following steps:

1. With host preparation, the VIB modules responsible for endpoint protection are installed on the hypervisor hosts in a vSphere cluster.

2. The integrated service is then deployed as a virtual appliance running partner services, also known as a service virtual machine (SVM).
3. The administrator configures/determines endpoint protection policies for the VMs.
4. The integrated service uses the endpoint protection API library to introspect and protect guest VMs from malware.

Because the endpoint protection enables SVMs to read and write specific files on the guest VMs, it provides an efficient way to optimize memory use and avoid resource bottlenecks.

New VMs: Automated Policy Enforcement

Endpoint protection policies are automatically applied and enforced when a new virtual machine is added. Automatic policy enforcement is enforced with the following steps:

1. A new ESXi host transport node is added to the vSphere cluster.
2. NSX-T (automatically) deploys the endpoint protection framework and SVMs.
3. A new guest VM is started on the host transport node and the security policies are applied.
4. The security policies will follow the VM when the VM migrates to another host transport node using vSphere vMotion.

The thin agent is installed as part of the full installation of VMware Tools. The default installation of VMware Tools does not include the thin agent. See Figure 3-9.




Figure 3-9. Automated policy enforcement on new VMs

Virus and Malware: Automated Quarantine with Security Tags

You can use endpoint protection policies with tags to automatically place virtual machines into quarantine when they are compromised. See Figure 3-10.

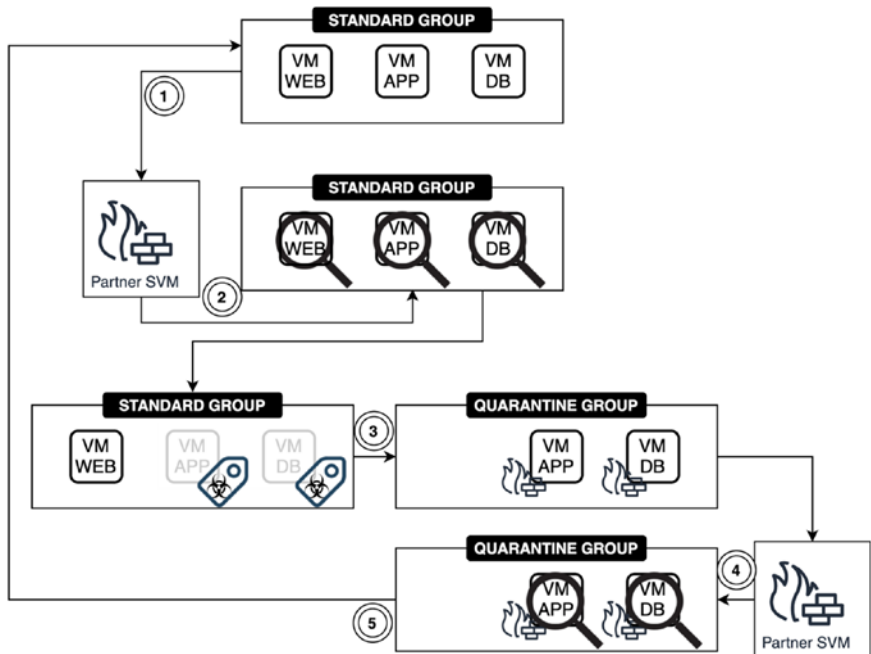


Figure 3-10. Automated quarantine using security tags

Here's how this works:

1. The three guest virtual machines operate as they normally would.
2. The antivirus SVM monitors activities on the guest virtual machines that are all in the STANDARD GROUP.
3. Malicious activities are detected and the infected virtual machines are assigned a tag. With the “quarantine” tag assignment, the infected virtual machines are now part of a security group called QUARANTINE GROUP. That group has a security policy bound to it in the form of a distributed firewall rule so that these infected virtual machines are not allowed to perform any form of network communication.
4. The antivirus SVM monitors activities on the guest virtual machines that are in the QUARANTINE GROUP.
5. When the virus has been removed, the “quarantine” tags are removed and the previously infected virtual machines are placed back into the STANDARD GROUP to operate normally again.

These steps happen automatically; no manual intervention is required.

Service Profile Creation for Endpoint Protection

You need to create a service profile for endpoint protection to define the level of protection you want to enforce for your virtual machines.

Service profiles enable you to choose protection levels for virtual machines by specifying the templates provided by the vendor.

A vendor can provide different levels of policies (red, orange, or green), which can be used to serve different types of virtual machines.

A red service profile provides complete antimalware to a PCI-type workload, and an orange service profile only provides basic antimalware protection to a regular virtual machine. See Figure 3-11.

<Add Screenshot>

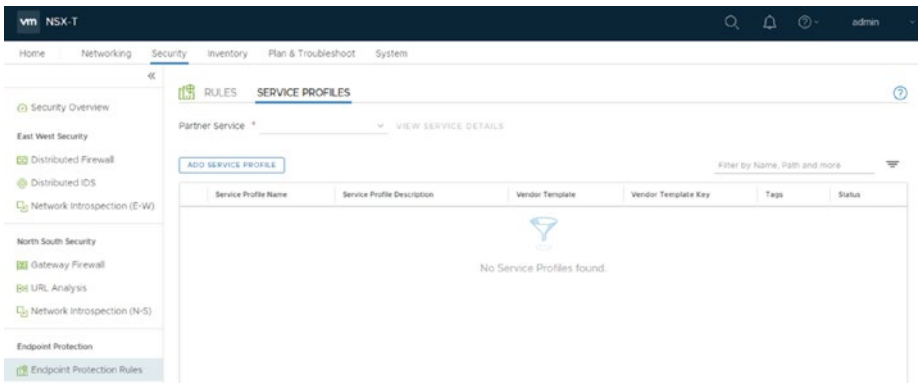


Figure 3-11. Service profile

Endpoint Protection Rules Configuration

When you create the service profile, you can specify an endpoint protection rule and apply it to a specific virtual machine or a group of virtual machines.

<Add Screenshot>

Guest Introspection Architecture

The guest introspection architecture is explained in this section. See Figure 3-12 and Table 3-2.

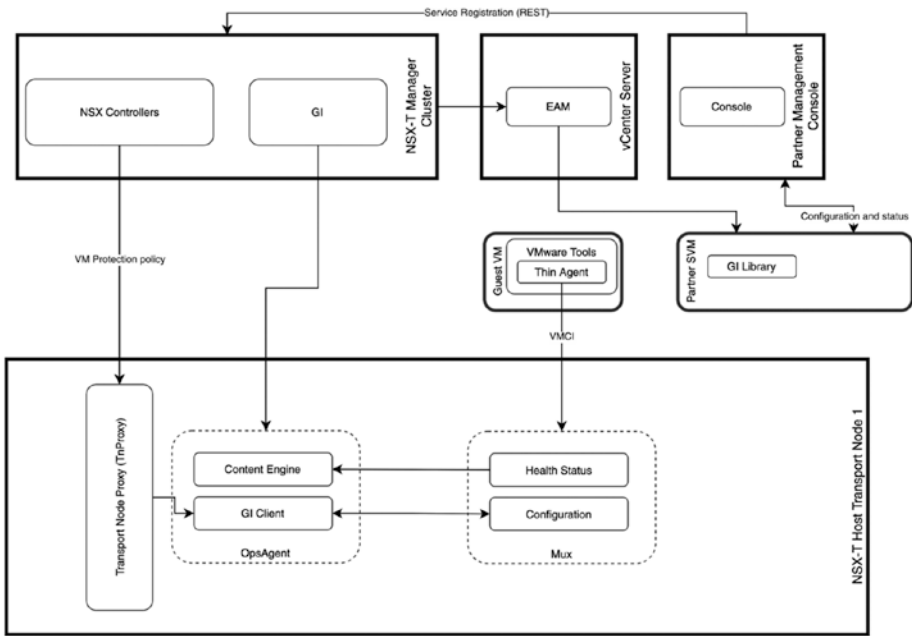


Figure 3-12. Guest introspection architecture

Table 3-2. Guest Introspection Architecture

Architecture	Description
Partner console	This is the web application provided by the security vendor to work with the guest introspection platform.
NSX Manager	This is the management plane appliance for NSX-T that exposes the API and graphical user interface to customers and partners for the configuration of Network and Security policies. For guest introspection, the NSX-T Manager also provides API and GUI to deploy and manage partner appliances.
Guest introspection SDK	This is the library that is provided by VMware so that the security vendor can consume this to integrate the service.

(continued)

Table 3-2. *(continued)*

Architecture	Description
Service VM	The security vendor provides a virtual machine that consumes the guest introspection SDK provided by VMware. It contains the intelligence to scan files or process events to detect virus or malware on the guest virtual machine. After the scanning has been performed, it sends a verdict or notification about the action taken by the guest virtual machine on the request.
Guest introspection host agent (context multiplexer)	This agent processes the configuration of endpoint protection policies. It also multiplexes and forwards messages from protected guest virtual machines to the service VM. It reports the health and status of the guest introspection platform and maintains records of the service VM configuration in the <code>muxconfig.xml</code> file.
Ops agent (context engine and guest introspection client)	This agent forwards the guest introspection configuration to the guest introspection host agent (context multiplexer). It also relays the health status of the partner solution to the NSX-T Manager.
EAM	The NSX-T Manager uses the ESXi agent manager to deploy a partner service VM on every host on the cluster that is configured for protection.
Thin agent	This agent is the file and/or network introspection agent running inside the guest virtual machine. It also intercepts file and network activities that are forwarded to the service VM through the host agent. This agent is part of VMware Tools. It replaces the traditional agent provided by antivirus or antimalware security vendors that are typically installed as applications/services manually. This thin agent is a generic and lightweight agent that facilitates offloading files and processes for scanning to the service VM provided by the vendor.

Summary

Now that you have finished reading this chapter, you should be able to identify the two different types of service insertion—network introspection service and endpoint protection. You also read about the difference between north-south and east-west service insertions. This chapter ended with an explanation of how endpoint protection works and some use cases, like automatic quarantining using tags of guest virtual machines.

The next chapter covers NSX-T network services, like Network Address Translation (NAT), Dynamic Host Configuration Protocol (DHCP), and Domain Name Services (DNS).

CHAPTER 4

NSX-T NAT, DHCP, and DNS Services

This chapter discusses NSX-T network services, including Network Address Translation (NAT), Dynamic Host Configuration Protocol (DHCP), and Domain Name Services (DNS).

It explains the differences between Source NAT, Destination NAT, and Reflexive NAT and explains how NAT64 works.

The DHCP section describes how and when DHCP is used and how to configure it.

Regarding DNS, the chapter covers the DNS capabilities that NSX-T offers and how to configure these.

Network Address Translation (NAT) Support

NAT can be configured on Tier-0 or Tier-1 gateways.

NSX-T supports three types of NAT options:

- Source NAT (SNAT)
- Destination NAT (DNAT)
- Reflexive NAT (stateless NAT)

SNAT and DNAT are both supported on Tier-0 or Tier-1 gateways when it operates in Active/Standby HA mode.

Reflexive NAT is only supported by a Tier-0 gateway when it operates in Active/Active HA mode.

Source Network Address Translation (SNAT)

Source NAT (SNAT) is used when you want to translate IP addresses that are configured inside your NSX-T domain to an IP address that is outside your NSX-T domain. A typical example of this is when you want to provide Internet access to internal virtual machines (that reside inside your NSX-T domain).

With an SNAT rule, the source IP address in the IP header of a packet changes. With this rule, you can also change (or translate) the source port in the TCP or UDP headers.

If you need to disable SNAT, you can do so using an “NO SNAT” rule.

In Figure 4-1, the WEB (172.16.10.0/24), APP (172.16.20.0/24), and DB (172.16.30.0/24) segments are translated with SNAT to a public IP address 100.100.10.100. That way, the virtual machines that are connected to the NSX-T segments can have Internet connectivity.

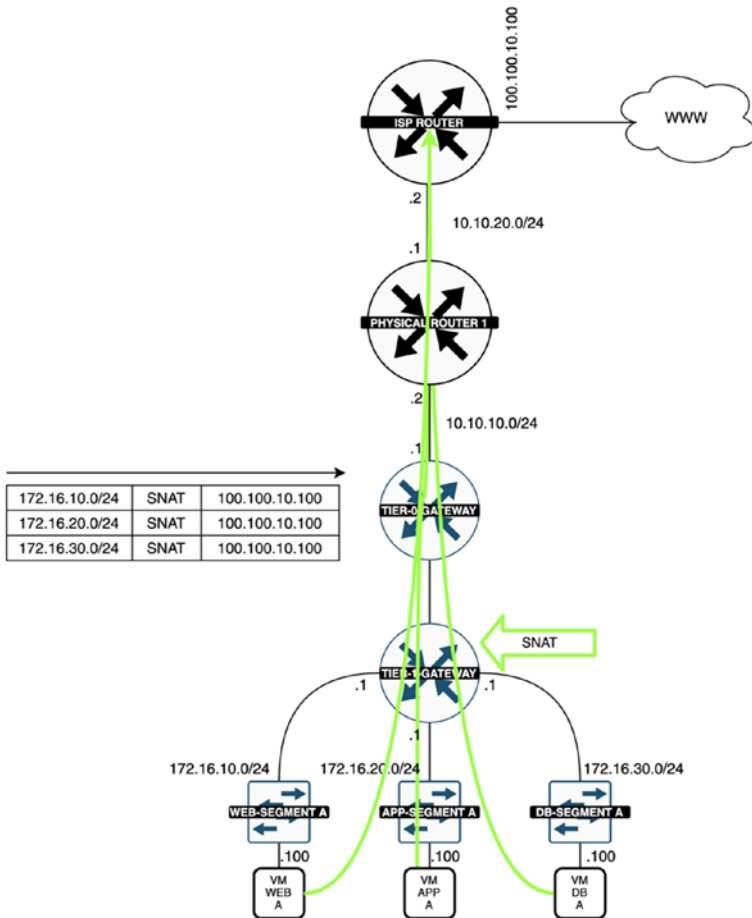


Figure 4-1. SNAT

Destination Network Address Translation (DNAT)

Destination NAT (DNAT) is used when you want to translate IP addresses that are configured outside your NSX-T domain to an IP address that is inside your NSX-T domain. A typical example of this is when you want to expose a web server that is configured on internal virtual machines (that reside inside your NSX-T domain) to the Internet so that users on the Internet can reach your web server.

With an DNAT rule, the destination IP address in the IP header of a packet changes. With this rule, you can also change (or translate) the destination port in the TCP or UDP headers.

If you need to disable DNAT, you can do this with a “NO DNAT” rule.

In Figure 4-2, the WEB server has the internal IP address of 172.16.10.100/24. The IP address 100.100.10.100 will be translated with DNAT to the private IP address 172.16.10.100. That way, Internet users can reach the web server inside your NSX-T domain.

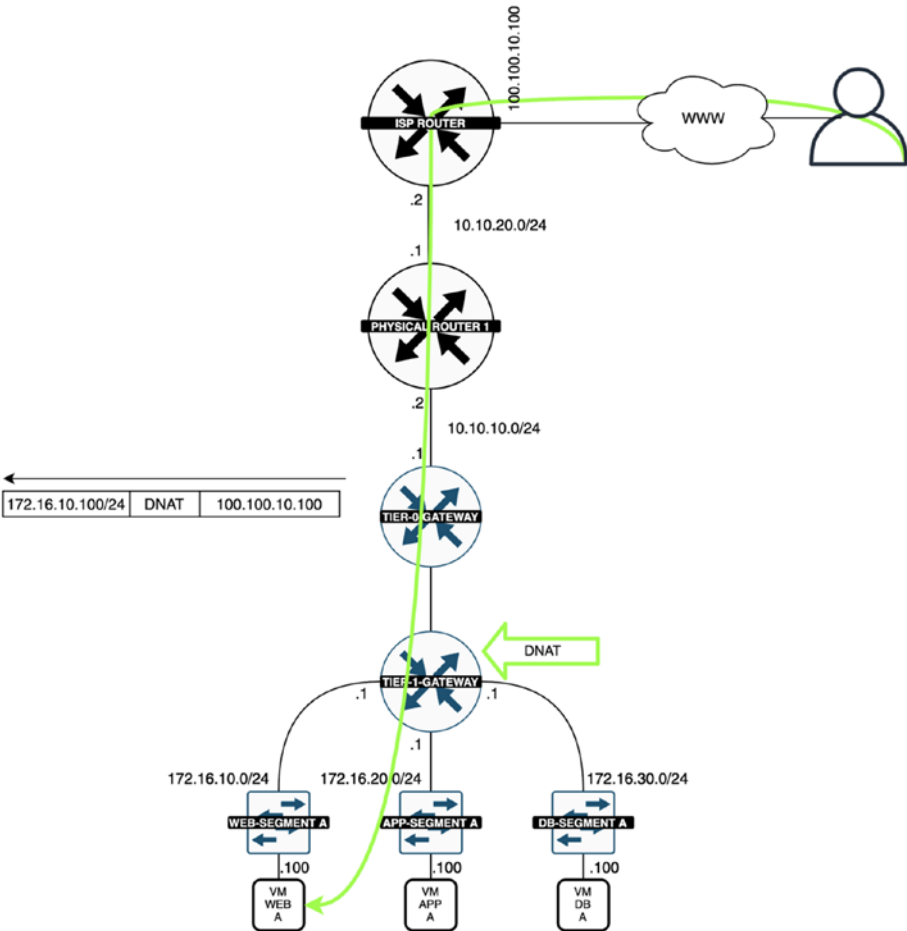


Figure 4-2. DNAT

Reflexive Network Address Translation (Reflexive NAT)

Reflexive NAT is configured when you have your Tier-0 gateway in Active/Active HA mode between two edge transport nodes. When you use SNAT or DNAT here (both stateful NAT protocols), this could lead to unexpected issues due to the asymmetric traffic paths.

Reflexive NAT is also called *stateless NAT*. You can configure a translation from a single source IP address to multiple destination IP addresses (range). You can also configure a range of source IP addresses, but keep in mind that you also have to configure a range of IP addresses to translate to.

In Figure 4-3, the source VM (172.16.10.100) on the inside network sends a packet to an outside client located on the Internet. The packet is first routed to the Tier-0 gateway that is hosted on NSX edge transport node 1. This creates a reflexive NAT entry: source IP 172.16.10.100 and translated IP 100.100.10.1. When the return traffic arrives (with the destination 100.100.10.100), the same reflexive NAT entry is used to translate 100.100.10.100 to 172.16.10.100.

When the other route is taken, NSX edge transport node 2 creates a reflexive NAT entry: source IP 172.16.10.100 and translated IP 100.100.10.200. When the return traffic arrives (with the destination 100.100.10.200), the same reflexive NAT entry is used to translate 100.100.10.200 to 172.16.10.100.

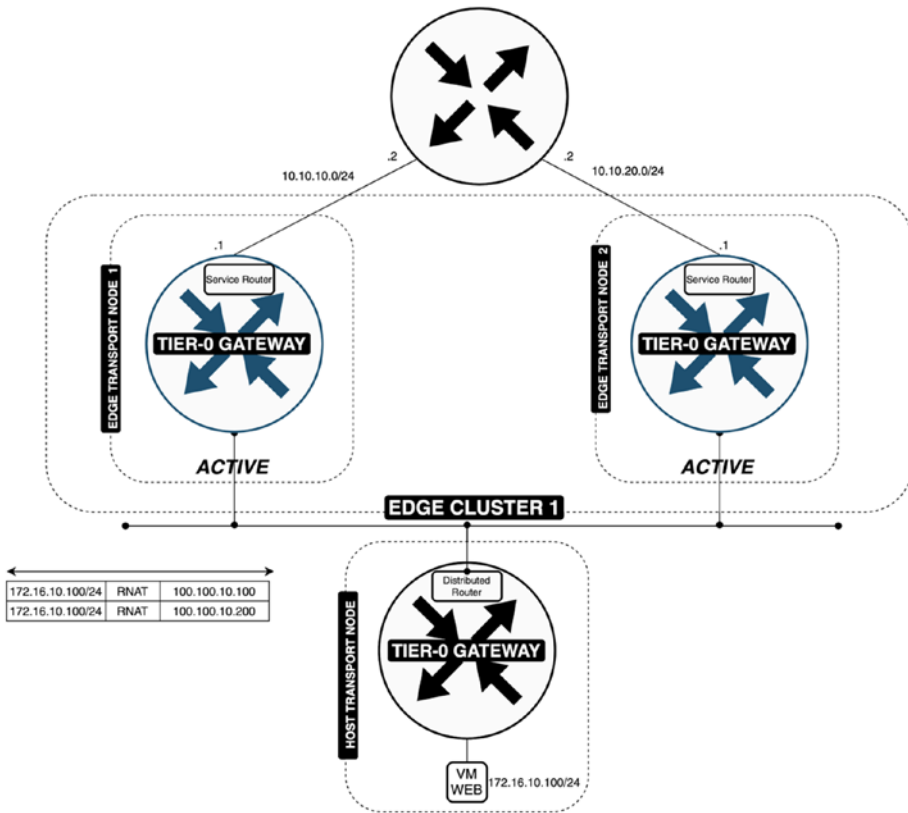


Figure 4-3. Reflexive NAT

SNAT and DNAT Configuration

SNAT or DNAT rules can be configured on Tier-0 and Tier-1 gateways. You can do this by navigating to Networking > Network Services > NAT > Select Gateway > Add NAT Rule, as shown in Figure 4-4.

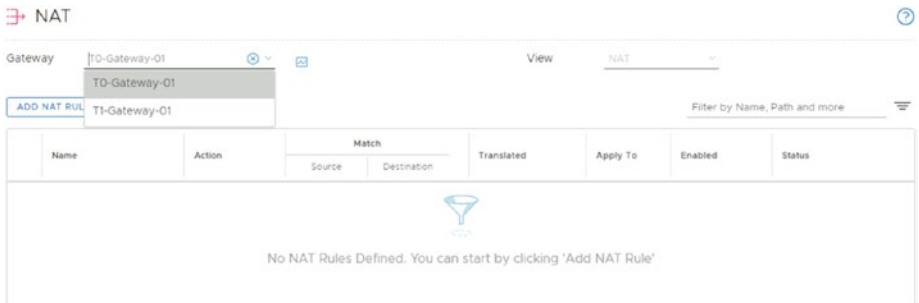


Figure 4-4. NAT configuring (start)

When you are configuring NAT, the parameters outlined in Table 4-1 can be used.

Table 4-1. NAT Configuration Parameters

Parameter	Description
Name	A name for the NAT rule.
Action	The action of the NAT rule if there is a match. The No SNAT/DNAT action disables the SNAT/DNAT rule and the original source/destination address in the IP packet is not translated.
Source IP	The source IP address or an IP address range in CIDR format. When you leave this text box blank, the NAT rule applies to all sources outside the local subnet (Any).
Destination IP	The destination IP address or an IP address range in CIDR format.
Translated IP	The new IP address as a result of NAT.
Service	Single service entry on which the NAT rule is applied.

(continued)

Table 4-1. (continued)

Parameter	Description
Firewall	<p>Match External Address: Matches the firewall rule with the external address of the NAT rule.</p> <p>Match Internal Address: Matches the firewall rule with the internal address of the NAT rule.</p> <p>Bypass: Skips the firewall stage.</p>
Applied To	<p>Select objects that this NAT rule applies to.</p> <p>Available objects are Tier-0 gateways, Tier-1 gateways, interfaces, labels, service instance endpoints, and virtual endpoints.</p>

For SNAT, you need to specify the source IP network and the translated IP address in the rule and use this rule for the OUTBOUND traffic direction.

For DNAT, you need to specify the destination IP network and the translated IP address in the rule and use this rule for the INBOUND traffic direction. See Figure 4-5.

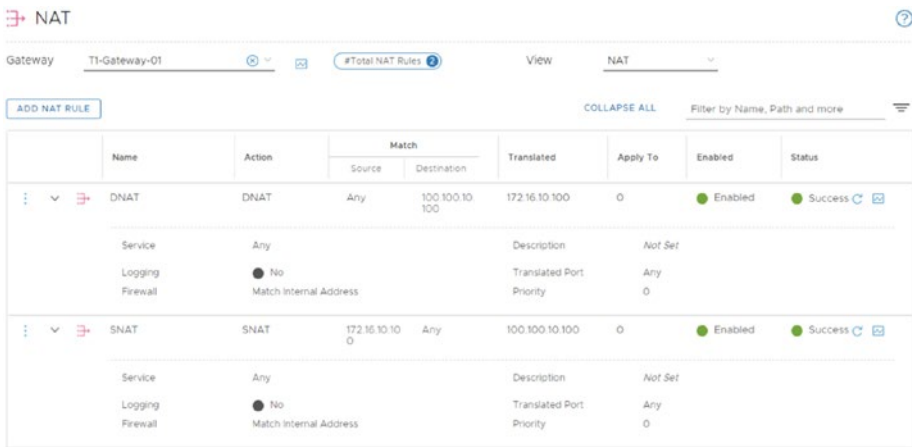


Figure 4-5. NAT configuration (finish)

No SNAT or DNAT Rule Configuration

To disable SNAT or DNAT, you can configure a NO SNAT or NO SNAT rule on the Tier-0 or Tier-1 gateway.

For NO SNAT, you need to specify the source IP network in the rule and use this rule for the OUTBOUND traffic direction.

For NO DNAT, you need to specify the destination IP network in the rule and use this rule for the INBOUND traffic direction. See Figure 4-6.

The screenshot shows the NAT configuration page for a Tier-1 Gateway (T1-Gateway-01). It displays two NAT rules:

	Name	Action	Match		Translated	Apply To	Enabled	Status
			Source	Destination				
⋮	DNAT	No DNAT	Any	100.100.10.100	Any	0	Enabled	Success
	Service	Any			Description	Not Set		
	Logging	<input checked="" type="radio"/> No			Translated Port	Any		
	Firewall	Match Internal Address			Priority	0		
⋮	SNAT	No SNAT	172.16.10.10/0	Any	Any	0	Enabled	In Progress
	Service	Any			Description	Not Set		
	Logging	<input checked="" type="radio"/> No			Translated Port	Any		
	Firewall	Match Internal Address			Priority	0		

Figure 4-6. NO NAT configuration (finish)

Reflexive NAT Configuration

When you want to apply NAT to a Tier-0 gateway and this gateway is configured in Active/Active HA mode across multiple edge transport nodes, the only option you can use is reflexive NAT (as stateful NATting is not possible in this design scenario).

The address translation is sequential, which means that the first address in the source range is translated to the first address in the translated range, and so on. Therefore, the two ranges (source and translated ranges) must be of equal size. See Figure 4-7.

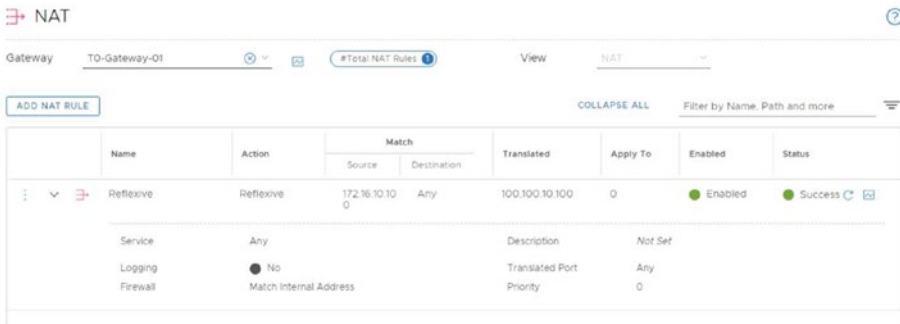


Figure 4-7. Reflexive NAT configuration

NAT Packet Flow

In this scenario, the Tier-1-A (Tenant A) performs SNAT. The private IP address 172.16.10.100 is translated to 100.100.10.100 in order to have the capability to be routed outside to the physical network.

When the private IP address 172.16.10.100 needs to reach the PROD VM (192.168.10.100), the NO NAT rule is used and the translation is disabled for this specific traffic route (when SNAT is configured on the gateway). See Figure 4-8.

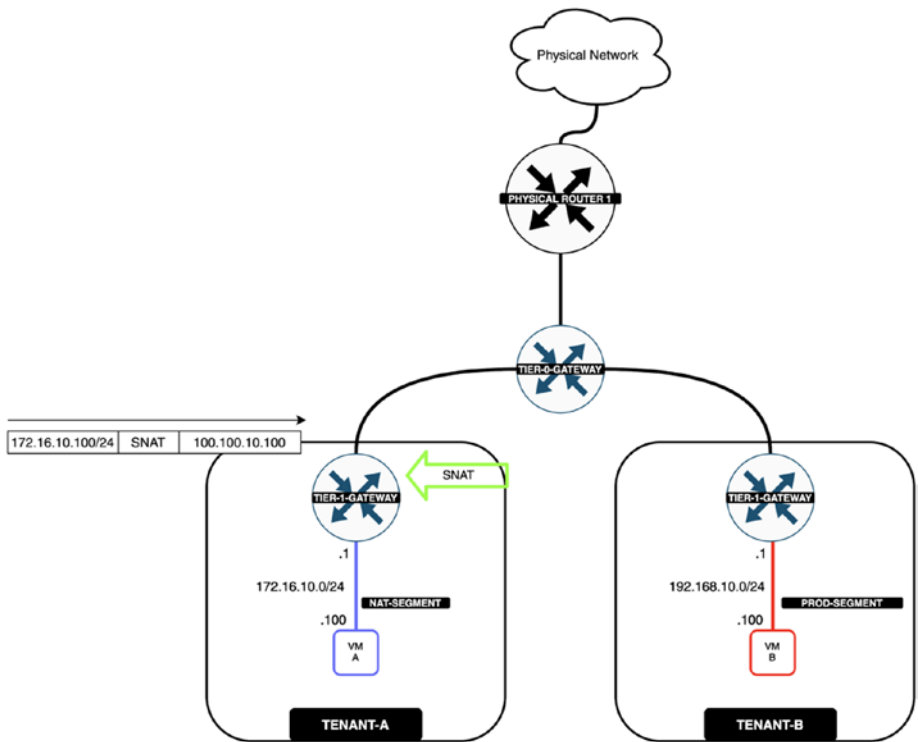


Figure 4-8. NAT packet flow

To prod (Figure 4-9):

The virtual machine (172.16.10.100) attached to NAT-SEGMENT sends a packet to its default gateway on T1-GW-A, which is the traffic with the source 172.16.10.100 and destination of the PROD VM (192.168.10.100).

The packet's destination (from the source 172.16.10.100) is the virtual machine (192.168.10.100), which is attached to PROD-SEGMENT.

1. The distributed router (DR) of T1-GW-A handles the packet.
2. The T1-GW-A DR makes the decision that this packet requires address translation, and this “NAT service” is provided by the service router (SR) of

the Tier-1 gateway on the edge transport node. T1-GW-A DR sends the packet to the local TEP interface.

3. The packet is then encapsulated with the Geneve header and is sent across the overlay tunnel. It arrives at the edge transport node (TEP interface).
4. The receiving TEP interface decapsulates the packet and sends the original packet to the SR of T1-GW-A for (NAT) processing. NAT (provided by the SR) translates the source address (172.16.10.100) to 100.100.10.100.
5. The packet with the translated IP (100.100.10.100) is sent to the DR of the T0-GW gateway for routing (edge transport node).
6. Based on its routing table (and destination), T0-GW routes the packet to the DR of T1-GW-B (edge transport node).
7. The T1-GW-B DR determines that the packet is destined for a remote VM and sends the packet to the local TEP interface (edge transport node).
8. (Edge) Transport Node B encapsulates the packet and sends it across the overlay tunnel (edge transport node).
9. The destination segment, PROD-SEGMENT, is attached to T1-GW-B on (host) transport node A. The DR of T1-GW-B routes the packet to the segment based on its routing table.
10. The packet arrives at the destination virtual machine (192.168.10.100).

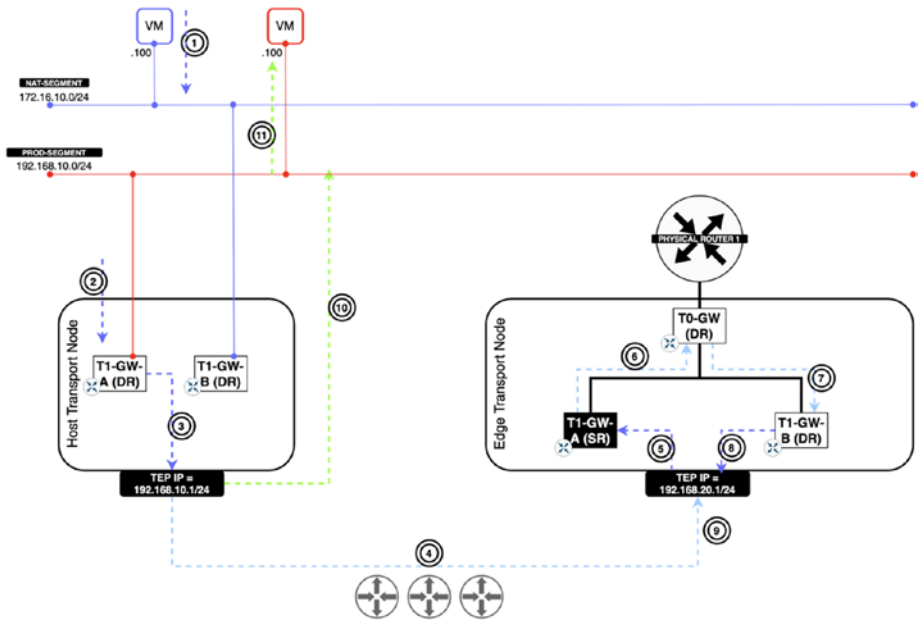


Figure 4-9. NAT packet flow in the production environment

To the Internet (Figure 4-10):

The virtual machine (172.16.10.100) attached to NAT-SEGMENT sends a packet to its default gateway on T1-GW-A.

The traffic with the source 172.16.10.100 and destination the Internet follows this flow:

1. The packet's destination (from the source 172.16.10.100) is the Internet.
2. The distributed router (DR) of T1-GW-A handles the packet.
3. The T1-GW-A DR makes the decision that this packet requires no address translation, and this "NO NAT service" is provided by the service router (SR) of the Tier-1 gateway on another (remote) host. The T1-GW-A DR sends the packet to the local TEP interface.

4. The packet is then encapsulated with the Geneve header and is sent across the overlay tunnel. It arrives at the remote host (TEP interface).
5. The receiving TEP interface decapsulates the packet and sends the original packet to the SR of T1-GW-A for (SNAT) processing. SNAT (provided by the SR) prevents translation of the source address (172.16.10.100).
6. The packet is sent to the DR of the T0-GW gateway for routing.
7. Based on its routing table (and destination), the SR of the T0-GW routes the packet to the physical router.
8. The packet destined for the Internet is routed to the physical network.

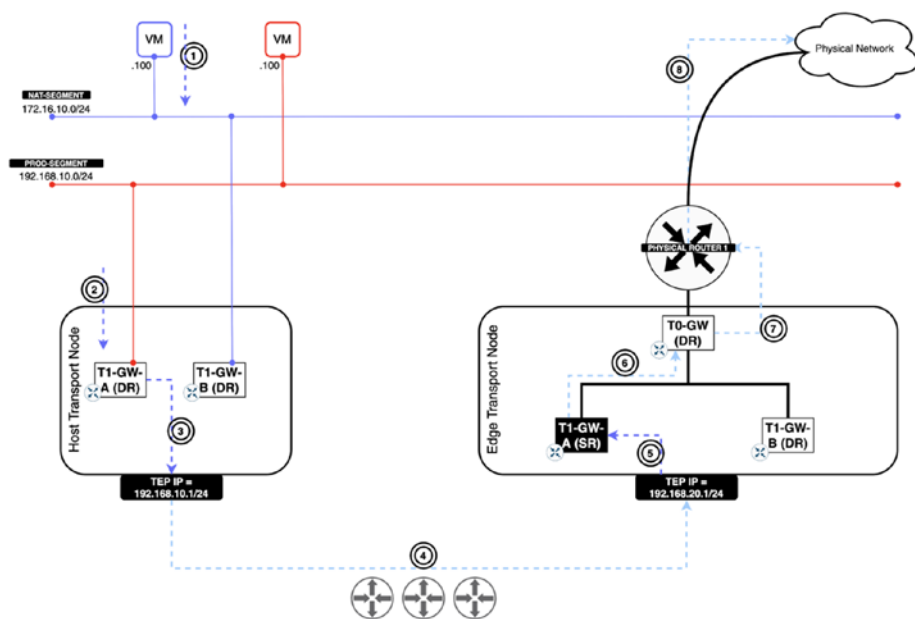


Figure 4-10. NAT packet flow to the Internet

NAT64

The NAT64 feature enables IPv6 to IPv4 connectivity. NAT64 is based on the RFC 6146 and RFC 6145 Internet Engineering Task Force (IETF) standards.

NAT64 allows an IPv6-only client to initiate communications to an IPv4-only server. IPv6 must initiate the traffic. See Figure 4-11.

NAT64 translates IPv6 packets to IPv4 packets and forwards them to the IPv4 network.

This functionality is designed so that you don't need to make changes to the IPv6 or IPv4 nodes.

At the time of this writing, NAT64 is only supported on Tier-0 gateways. NAT64 is a stateful service and requires the Tier-0 gateway to be deployed in Active/Standby HA mode.

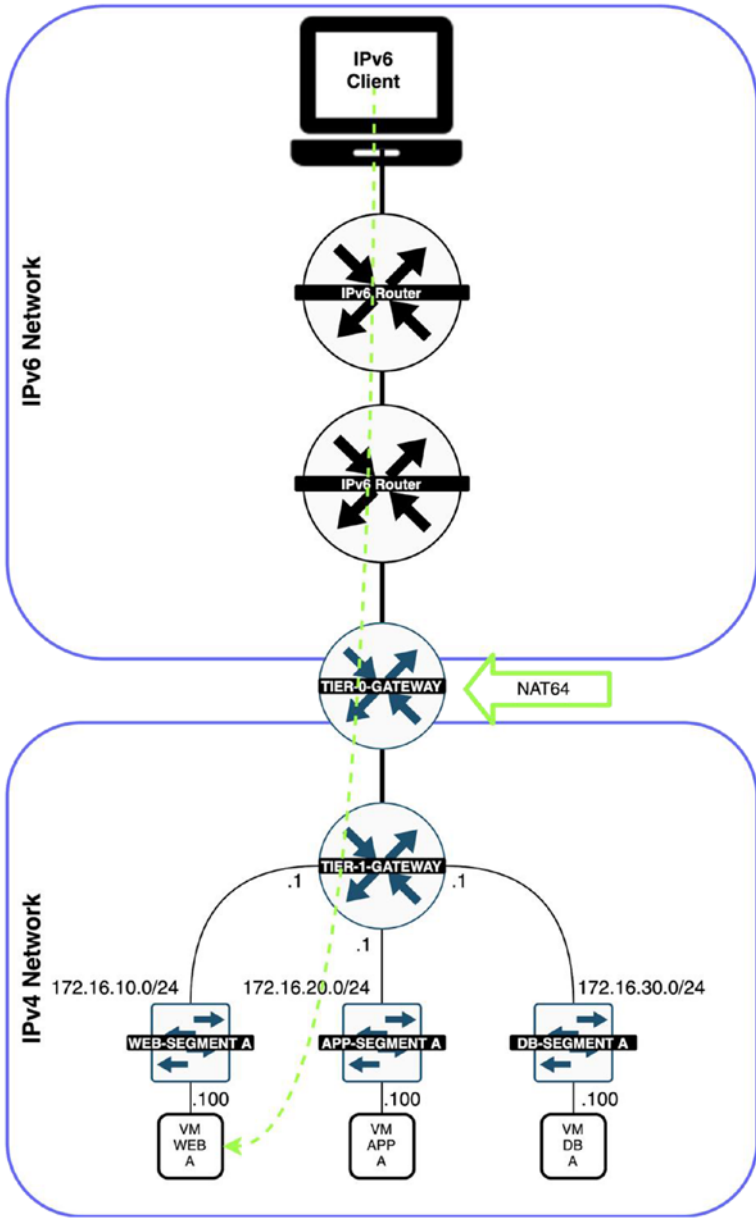


Figure 4-11. NAT64

NAT64 Use Cases

NAT64 is useful when you have to merge two different networks of two different companies, due to a merger for example. Another use case is if your (new) IPv6 servers need to access (legacy) IPv4 network servers (services).

NAT64 Requirements

Before you configure NAT64 on the Tier-0 gateway, you must meet the following requirements (Figure 4-12).

- The Tier-0 gateway must be configured with at least two interfaces:
 - One uplink interface that resides in the IPv6 network
 - One downlink interface that resides in the IPv4 network
- An IPv4 IP address pool that can consist of one single IPv4 address or multiple IPv4 addresses (that is required to translate the IPv6 source IP addresses).

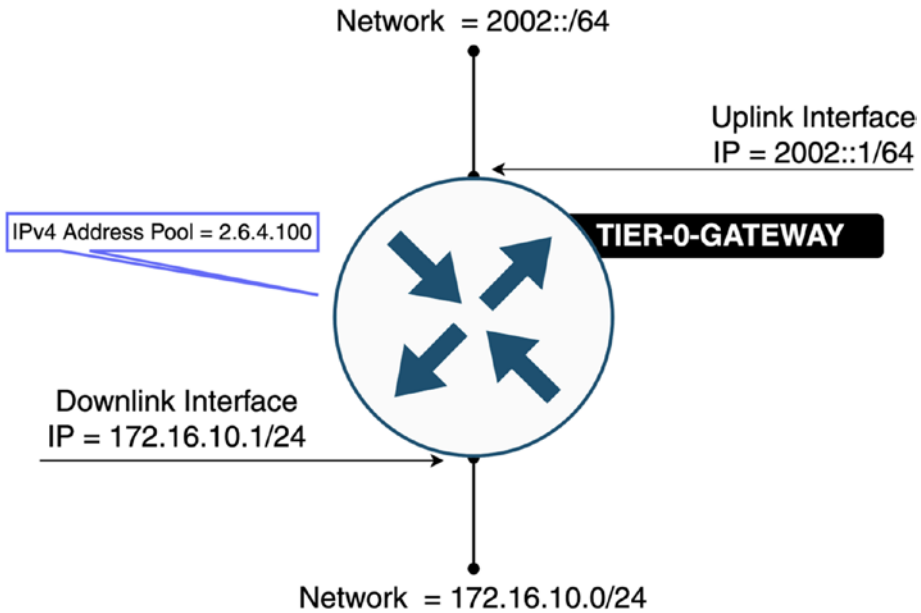


Figure 4-12. NAT64 interface names

NAT64 Limitations

At the time of this writing, NSX-T has the following limitations in regards to the NAT64 feature:

- Only TCP and UDP are supported. Other protocol-type packets are discarded (including ICMP).
- Path MTU Discovery (PMTUD) is not supported.
- Broadcast or multicast packets are not supported.
- NAT64 rules can be applied only to interfaces or tags. Applying rules to groups is not supported.

NAT64 Tables

NAT64 maintains two tables to translate IPv6 packets to IPv4 packets:

- A Binding Information Base (BIB) table, which maintains the IPv6 to IPv4 mappings. See Figure 4-13.
- A Session table, which maintains the currently established TCP/UDP sessions. See Figure 4-14.

Incoming IPv6 TCP/UDP Packet	Translated IPv4 TCP/UDP Packet
IPv6 SRC IP + IPv6 SRC Port	IPv4 SRC IP + IPv4 SRC Port

Figure 4-13. BIB table

Source IPv6 Transport Address	Destination IPv6 Transport Address	Source IPv4 Transport Address	Destination IPv4 Transport Address	Session Timer
IPv6 SRC IP + IPv6 SRC Port	IPv6 SRC IP + IPv6 DST Port	IPv4 SRC IP + IPv4 SRC Port	IPv4 SRC IP + IPv4 DST Port	Maximum session lifetime

Figure 4-14. Session table

NAT64 Packet Flow

The following steps explain the NAT64 packet flow process (Figure 4-15):

1. The IPv6 client sends traffic to the IPv4 virtual machine.
2. For the incoming packet from the IPv6 network (2000::2, 2004::172.16.10.100), the Tier-0 gateway constructs the IPv4 packet based on the NAT64 rule. It derives the IPv4 destination address (172.16.10.100) from the IPv6 packet version 4 portion of its destination IP (2004::172.16.10.100) and assigns the source address from the IPv4 address pool (2.4.6.100).

3. The IPv4 SRC port is either the same as the IPv6 SRC Port or a free port. In the example, the IPv6 SRC port is TCP/1000 and the IPv4 SRC port is TCP/50000.
4. The IPv4 DST port is the same as the IPv6 DST port: port TCP/22 in Figure 4-15.
5. NAT64 copies the remaining fields from IPv6 to IPv4, such as TTL, TOS, and payload.
6. After the IPv4 packet (2.6.4.100, 172.16.10.100) is constructed, it is delivered to the IPv4 pipeline, which forwards the packet based on its IPv4 destination address.

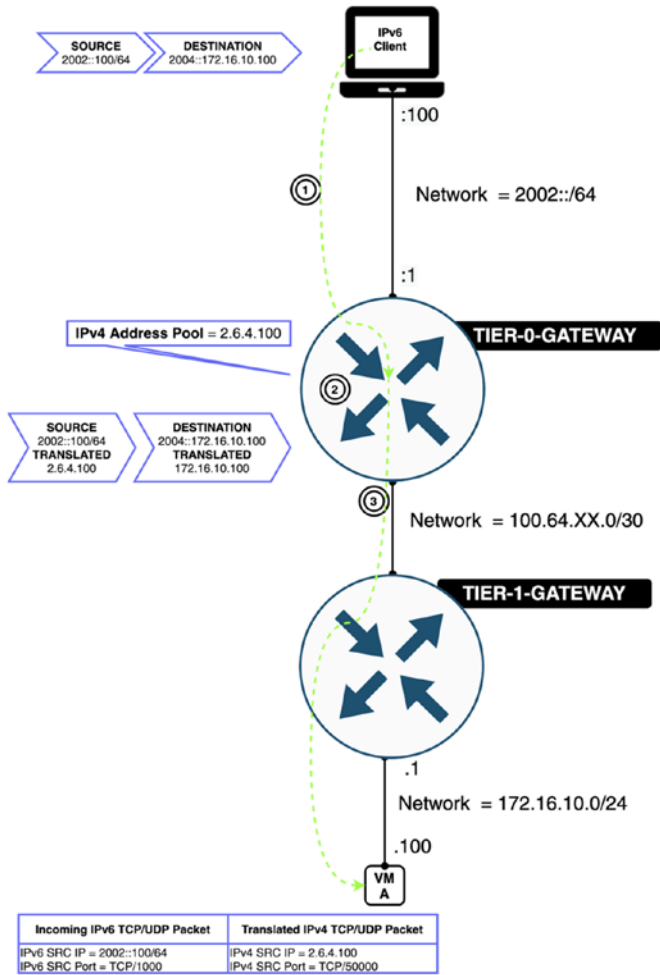


Figure 4-15. NAT64 packet flow

Source IPv6 Transport Address	Destination IPv6 Transport Address	Source IPv4 Transport Address	Destination IPv4 Transport Address	Session Timer
IPv6 SRC IP = 2002::100/64	IPv6 SRC IP = 2004::172.16.10.100	IPv4 SRC IP = 2.6.4.100	IPv4 SRC IP = 172.16.10.100	2 hours
IPv6 SRC Port = TCP/1000	IPv6 DST Port = TCP/22	IPv4 SRC Port = TCP/50000	IPv4 DST Port = TCP/22	

NAT64 Rule Configuration

NAT64 rules can be configured on the Tier-0 gateway. You can do this by navigating to Networking ► Network Services ► Select Gateway ► Select View: NAT64 ► Add NAT Rule.

When you are configuring NAT64, the following parameters outlined in Table 4-2 can be used.

Table 4-2. NAT64 Configuration Parameters

Parameter	Description
Name	A name for the NAT64 rule.
Action	The only supported action is NAT64, which translates IPv6 addresses to IPv4.
Source	An optional field. IPv6 or IPv6 address range in CIDR format.
Destination	A required field. IPv6 or IPv6 address range in CIDR format. The prefix must be /96 because the destination IPv4 IP is embedded as the last four bytes in the IPv6 address.
Translated	This address is the IPv4 address or address pool for the source IPv6 address.
Applied To	NAT64 rules can only be applied to uplink interfaces or tags.

You can optionally specify the matching service and translated port too.

DHCP Services

Dynamic Host Configuration Protocol (DHCP) allows network clients to obtain network configuration settings—including the IP address, subnet mask, default gateway, DNS servers, and much more—automatically from

a DHCP server. The DHCP process (or service) that runs on a server is typically called a *DHCP server*.

A DHCP server is either placed on a specific network segment so that it can cater to the network clients of that same segment, or you can create a DHCP relay on a router/gateway so that the DHCP requests are related to a DHCP server that is residing on a different network.

Never associate a DHCP server and DHCP relay service to the same segment, as this will introduce conflicts.

DHCP Architecture

A DHCP server configuration can either be applied to a Tier-1 or Tier-0 gateway or to a specific segment.

Using NSX-T, the DHCP services are configured on Tier-1 or Tier-0 gateways with the following rules:

- A gateway always has a distributed router component (DR).
- A gateway has one or more service router components (SR) when it is a Tier-0 gateway or when it is a Tier-1 gateway with routing services enabled, such as NAT or DHCP for example.
- An NSX edge cluster must be created to configure DHCP on a Tier-1 or Tier-0 gateway.
- A DHCP server and DHCP relay server can be configured on Tier-1 or Tier-0 gateway. See Figure 4-16.

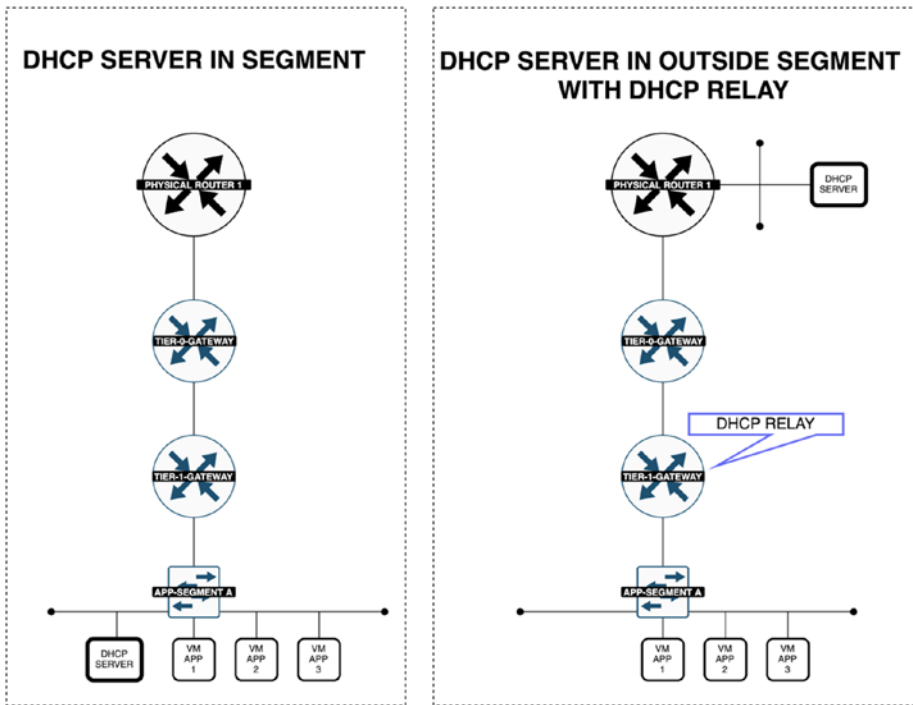


Figure 4-16. DHCP server location

DHCP Use Cases

The DHCP service can run as a server on an NSX-T Edge node. Typical use cases for DHCP (briefly discussed in the previous chapter) are as follows:

- Automate IP address assignment to network clients, including IP address, subnet mask, default gateway IP address, DNS IP address, and much more (optional) network parameters that you require.
- A “logical” DHCP server can be configured for a specific segment and pool.

- When you have a virtual machine with a virtual Network Interface Card (vNIC) that is running a DHCP client, this client acquires network parameters dynamically and configures this on the vNIC.
- You can create a DHCP cluster with multiple edge transport nodes. The (logical) configured DHCP servers are then highly available following the edge cluster guidelines.

DHCP Workflow

To create a DHCP server, you need to follow these steps:

1. Create a DHCP server profile.
2. Configure a Tier-0 or Tier-1 gateway with the DHCP server profile.
3. Create a segment and configure the DHCP IP range.
4. Attach the virtual machine vNIC to the segment.
5. Configure the virtual machine operating system so that the vNIC can acquire its network settings using DHCP.

DHCP Server Profile Creation

The DHCP server profile must be created and attached to a Tier-0 or Tier-1 gateway to provide the DHCP IP address to the virtual machine vNICs.

To create a DHCP server profile, navigate to Networking ► DHCP ► Add DHCP Profile, as shown in Figure 4-17.

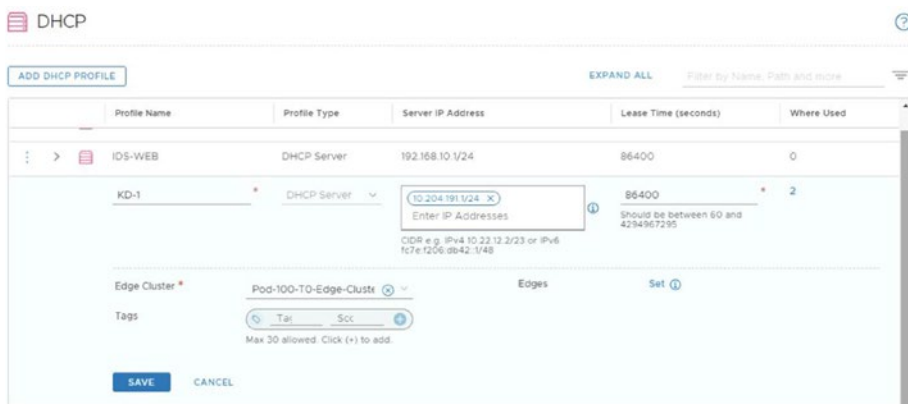


Figure 4-17. DHCP profile configuration

DHCP Server on a Tier-0/Tier-1 Gateway

To configure the DHCP server on the Tier-0/Tier-1 gateway, you edit the IP address management by changing the type of server using Set DHCP Configuration.

The configuration for the DHCP on the Tier-0 or Tier-1 gateway is similar. When you edit the gateway, you will see a DHCP section, as shown in Figure 4-18. Click the Set DHCP Configuration option.

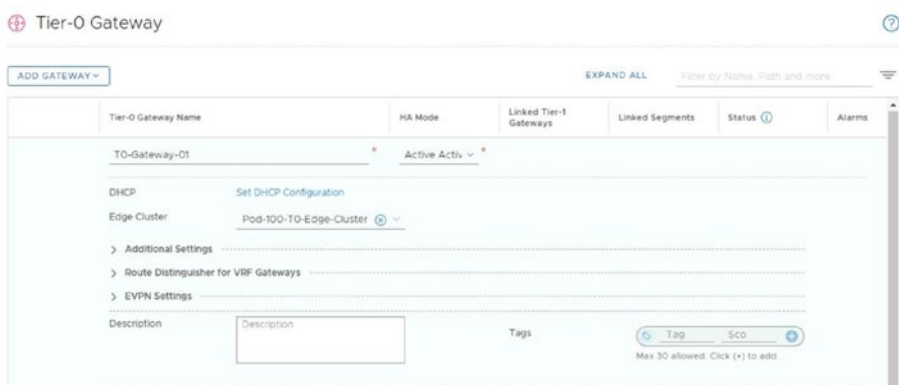


Figure 4-18. DHCP server configuration on a Tier-0/Tier-1 gateway (1)

By default, the DHCP configuration is set to No Dynamic IP address Allocation, as shown in Figure 4-19.

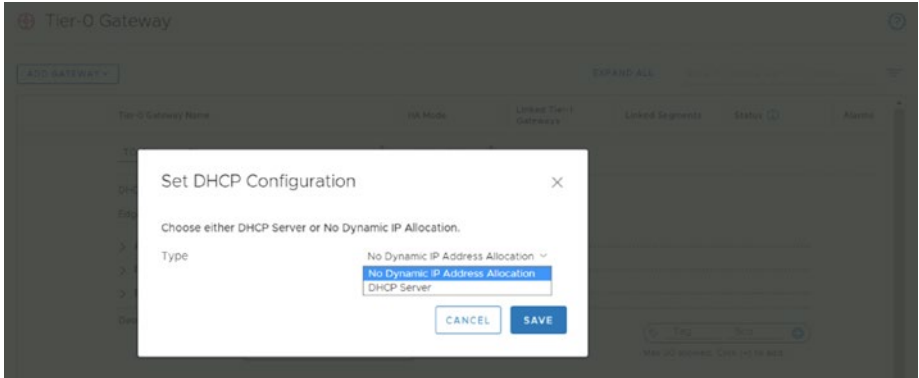


Figure 4-19. DHCP server configuration on a Tier-0/Tier-1 gateway (2)

Select DHCP Server and then choose a DHCP Server Profile, as shown in Figure 4-20.

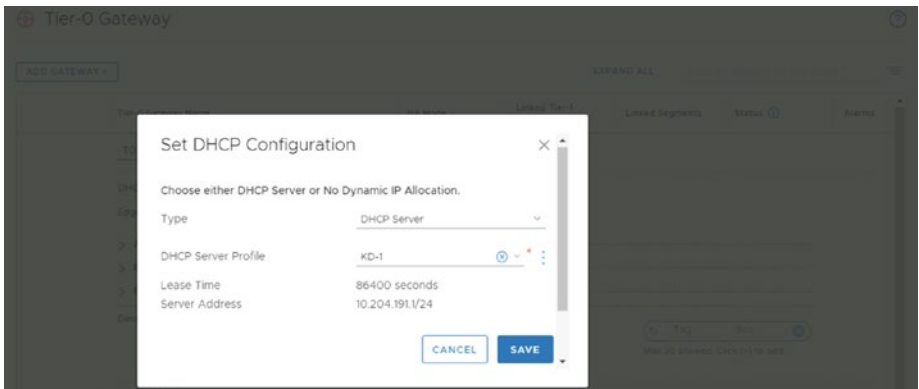


Figure 4-20. DHCP server configuration on a Tier-0/Tier-1 gateway (3)

When you click Save, you see that 1 Server is configured. See Figure 4-21.

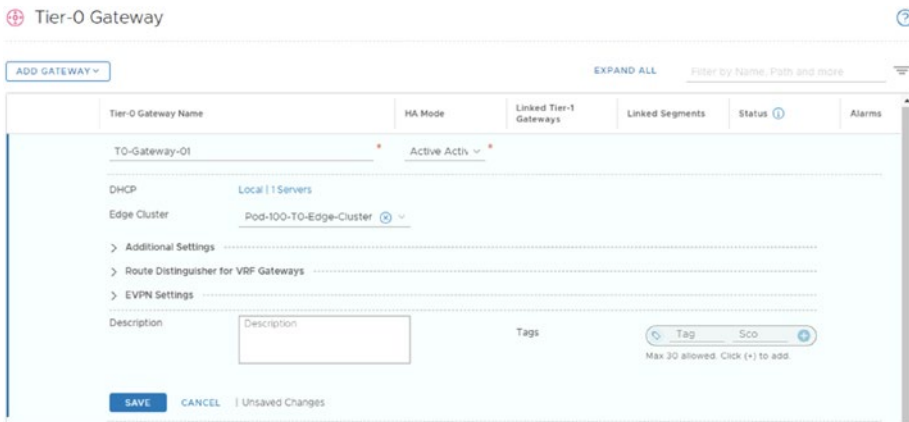


Figure 4-21. DHCP Server configuration on a Tier-0/Tier-1 gateway (4)

DHCP Configuration on a Segment

The DHCP server can be configured on a specific segment. To configure the DHCP server on a segment, select Edit DHCP Config from the segment, as shown in Figures 4-22 through 4-24.

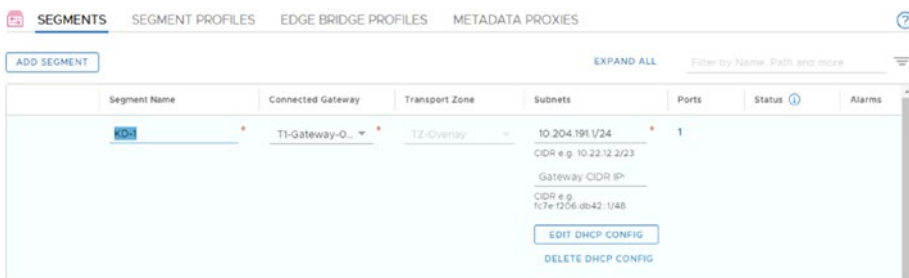


Figure 4-22. DHCP server configuration on a segment (1)

Set DHCP Config

Segment KD-1

IPv4 Gateway 10.204.191.1/24 #DHCP Ranges ⓘ IPv6 Gateway Not Set #DHCP Ranges ⓘ

DHCP Type * Local DHCP Server ⓘ DHCP Profile * KD-1 ⓘ

IPv4 Server IPv6 Server

Settings | Options

DHCP Config Enabled ⓘ

DHCP Server Address * 10.204.191.2/24
CIDR e.g. 10.22.12.2/23

DHCP Ranges 99 Maximum | Format 172.16.14.10-172.16.14.100 or 172.16.14.0/24 | Please verify that IP addresses in this range are not in use prior to modifying the DHCP range to avoid duplicate IP address allocation

10.204.191.200-10.204.191.225 X
Enter DHCP Ranges

Lease Time (seconds) 86400

DNS Servers 10.203.0.5 X
Enter IP Addresses
e.g. 10.10.10.10

Figure 4-23. DHCP server configuration on a segment (2)

Set DHCP Config

Segment KD-1

IPv4 Gateway 10.204.191.1/24 #DHCP Ranges ⓘ IPv6 Gateway Not Set #DHCP Ranges ⓘ

DHCP Type * Local DHCP Server ⓘ DHCP Profile * KD-1 ⓘ

IPv4 Server IPv6 Server

Settings | Options

DHCP Config Enabled ⓘ

DHCP Server Address * 10.204.191.2/24
CIDR e.g. 10.22.12.2/23

DHCP Ranges 99 Maximum | Format 172.16.14.10-172.16.14.100 or 172.16.14.0/24 | Please verify that IP addresses in this range are not in use prior to modifying the DHCP range to avoid duplicate IP address allocation

10.204.191.200-10.204.191.225 X
Enter DHCP Ranges

Lease Time (seconds) 86400

DNS Servers 10.203.0.5 X
Enter IP Addresses
e.g. 10.10.10.10

- Gateway DHCP Server - Central DHCP service that dynamically assigns IP to the VMs on all the segments that are connected to the gateway and using Gateway DHCP
- Local DHCP Server - Defines DHCP server that is local to the segment and not available to the other segments in the network. A local DHCP server provides a dynamic IP assignment service only to the VMs that are attached to the segment
- DHCP Relay - Defines DHCP relay service that is local to the segment and not available to the other segments in the network. The DHCP relay service relays the DHCP requests of the VMs that are attached to the segment to the remote DHCP servers

Figure 4-24. DHCP server configuration on a segment (3)

DHCP Server Status Verification

You can check the DHCP server status (where this DHCP profile is used) from the DHCP page, as shown in Figure 4-25.

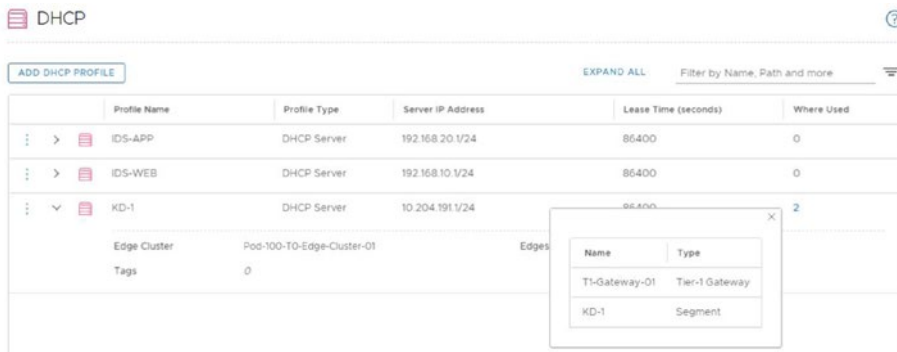


Figure 4-25. DHCP server verification

DHCP Relay Profile Creation

On the DHCP configuration page of the NSX UI, you can create DHCP servers to handle DHCP requests. You can also create DHCP relay services to relay the DHCP traffic to the external DHCP servers.

To add a DHCP relay, select DHCP Relay from the Profile Type drop-down menu and enter one or more IP addresses for the service of your actual DHCP servers. See Figure 4-26.

Profile Name	Profile Type	Server IP Address	Lease Time (seconds)	Where Used
RELAY	DHCP Relay	10.203.0.5	0	

Tags: Tag: Sco (Max 30 allowed. Click (+) to add.)

Buttons: SAVE, CANCEL

Figure 4-26. DHCP profile configuration (1)

Make sure you select the appropriate profile in the Segment settings, as shown in Figures 4-27 and 4-28.

Set DHCP Config

Segment: KD-2

IPv4 Gateway: 10.204.192.1/24 (#DHCP Ranges)

IPv6 Gateway: Not Set (#DHCP Ranges)

DHCP Type: DHCP Relay

DHCP Profile: RELAY (servers: 10.203.0.5)

Warning: IPv4 Server & IPv6 server settings are not supported for Relay

IPv4 Server: IPv6 Server

Settings | Options

DHCP Config: Disabled

DHCP Server Address: Enter CIDR (CIDR e.g. 10.22.12.2/23)

DHCP Ranges: 99 Maximum | Format 172.16.14.10-172.16.14.100 or 172.16.14.0/24 | Please verify that IP addresses in this range are not in use prior to modifying the DHCP range to avoid duplicate IP address allocation

Lease Time (seconds): Default value is 86400

DNS Servers: Enter IP Addresses (e.g. 10.10.10.10)

Buttons: CANCEL, APPLY

Figure 4-27. DHCP profile configuration (2)



Figure 4-28. DHCP profile configuration (3)

DHCP Relay Server on a Tier-0/Tier-1 Gateway Configuration

A DHCP relay server is configured through the DHCP profile and this profile is selected in the Segment settings.

DNS Services

To translate hostnames to IP addresses and the other way around, you use a DNS service. This DNS service typically runs on a (DNS) server and this server responds to queries (questions) that come from a DNS client.

NSX-T does not provide this DNS service, but it does support a DNS relay that acts as a caching and forwarding server. See Figure 4-29.

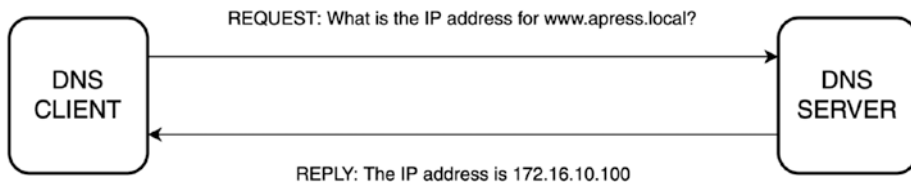


Figure 4-29. DNS client/server communication

DNS Forwarder

NSX-T will forward the DNS client requests to an external DNS server using its DNS forwarding feature.

This DNS forwarder is a server that forwards all DNS requests coming from DNS clients to the external DNS servers.

When a request is received (that has previously been made by another DNS client), the DNS forwarder can respond to it without sending another request to the external DNS server, unless the cached entry (TTL) has expired.

The Tier-0 and Tier-1 gateways can be configured with this DNS forwarder feature.

On Tier-0 and Tier-1 gateways, the DNS forwarder acts as a DNS relay to forward DNS client requests to external DNS servers. See Figure 4-30.

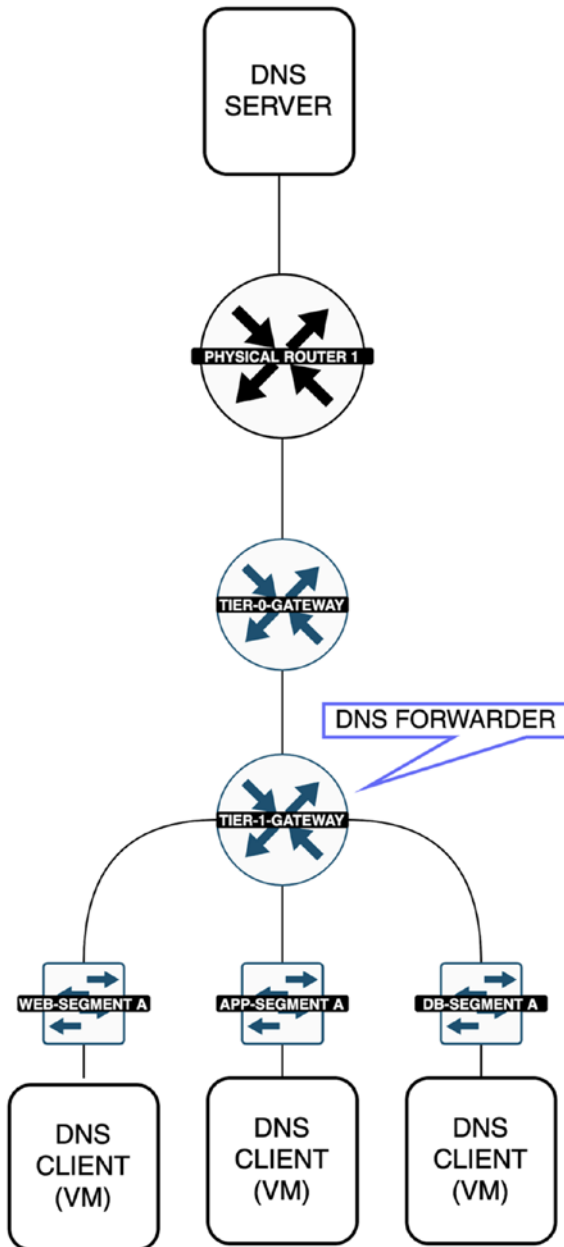


Figure 4-30. DNS client/server communication with a DNS forwarder

DNS Forwarder Benefits

The DNS forwarder feature uses the cached (DNS record) entry that is in the local caching database to respond to the DNS query requests. This reduces the load on the external DNS servers and improves the performance.

It is also possible to hide internal DNS information from external networks by delegating the internal DNS resolution to the internal DNS servers.

DNS Services and DNS Zones Configuration

You can configure a DNS forwarder by navigating to Networking ► IP Management ► DNS.

You first need to configure a (default) DNS zone using the DNS Zones tab. See Figure 4-31.

When you are configuring a DNS zone, the parameters outlined in Table 4-3 can be used.

Table 4-3. DNS Configuration Parameters

Parameter	Description
Zone name	A name for the DNS forwarder.
Domain	The default value is ANY and is the default DNS forwarder for all the domains except for those mentioned in the new fully qualified domain name (FQDN) zone.
DNS servers	The upstream DNS server IP address.
Source IP	The IP address that the DNS forwarder uses as the source IP address to send the DNS query to external DNS servers and receive the DNS reply from external DNS servers.
Description	An optional description for the DNS forwarder.
Tags	Optional tags for the DNS forwarder.

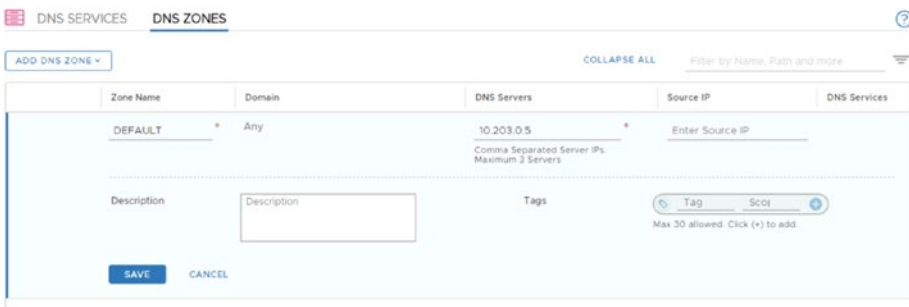


Figure 4-31. DNS zone configuration

Then you need to configure the DNS service using the DNS Services tab. When you configure the DNS service, the parameters outlined in Table 4-4 can be used. See Figure 4-32.

Table 4-4. DNS Service Configuration Parameters

Parameter	Description
Name	A name for the DNS service.
Tier-0/Tier-1 gateway	The gateway where you want to apply the DNS service. The Tier-0 gateway is only supported when configured in Active/Standby mode.
DNS service IP	This IP address is the listening IP address for DNS requests.
Default zone	The DNS client requests are forwarded to the default zone unless configured to relay the request to the conditional zone.

The screenshot displays the 'DNS SERVICES' configuration page in NSX-T. At the top, there's a breadcrumb 'DNS SERVICES' and 'DNS ZONES'. Below that, a button 'ADD DNS SERVICE' is visible. A table lists existing services, with 'DNS-FORWARD' selected. Below the table, the configuration form for 'DNS-FORWARD' is shown. It includes fields for 'Description', 'FQDN Zones' (with a dropdown and 'Maximum 5 FQDN Zones' note), 'Log Level' (set to 'Info'), and 'Admin Status' (set to 'Enabled'). There are also 'Tags' and a 'SAVE' button.

Figure 4-32. DNS services configuration

DNS Forwarder Verification

Once you have configured the DNS forwarder, you can verify the service status on an NSX-T Edge transport node using CLI (through the console or using an SSH connection).

Run the `get dns-forwarder status` command to query the DNS forwarder service running on NSX-T Edge transport node.

Run the `get dns-forwarder config` command to query the DNS forwarder configuration on the NSX-T Edge transport node.

The high availability for the DNS forwarder is active-standby, depending on the NSX-T Edge cluster.

Summary

This chapter explained what Network Address Translation (NAT) is and what the different NAT variations are, including Source NAT (SNAT), Destination NAT (DNAT), Reflexive NAT, and NAT64. It discussed specific use cases in which these are used.

You learned about NSX-T's current DHCP and DNS services, and how you can configure them.

The next chapter discusses the NSX-T load balancer's architecture and components.

CHAPTER 5

NSX-T Load Balancing

This chapter explains the NSX-T load balancer architecture and components.

The chapter identifies the load balancer deployment modes and explains how to create and configure a load balancer by configuring the Layer-4 and Layer-7 virtual servers and server pools.

NSX-T Load Balancing Use Cases

The NSX-T load balancer helps distribute incoming requests across multiple (virtual machine) servers. This not only offers high availability, but also allows better performance/throughput for applications.

You typically use load balancing:

- When you want to prevent server/service outages.
- When you require faster response times because the (client request) load is spread across multiple servers.
- When you want to steer the load toward a specific server in the pool.
- When a server in the pool needs maintenance. You can stop traffic from reaching that server and you can use other servers in the pool to guarantee the uptime of your application.

Layer-4 Load Balancing

The Layer-4 load balancer is a connection-based load balancer that supports the UDP and TCP protocols.

Figure 5-1 shows a connection coming in from the client that listens on TCP port 443. Based on this port and the configured server pools and profiles, the request will be forwarded to one of the servers in the server pool.




Figure 5-1. Layer-4 load balancing

Layer-7 Load Balancing

The Layer-7 load balancer is a content-based load balancer that supports HTTP and HTTPS. This load balancer can also enable URL manipulation through user-defined rules.

Figure 5-2 shows a client making two requests to two different DNS FQDN names. Based on the incoming FQDN, the Layer-7 load balancer makes a decision to send the request to a particular pool of servers.

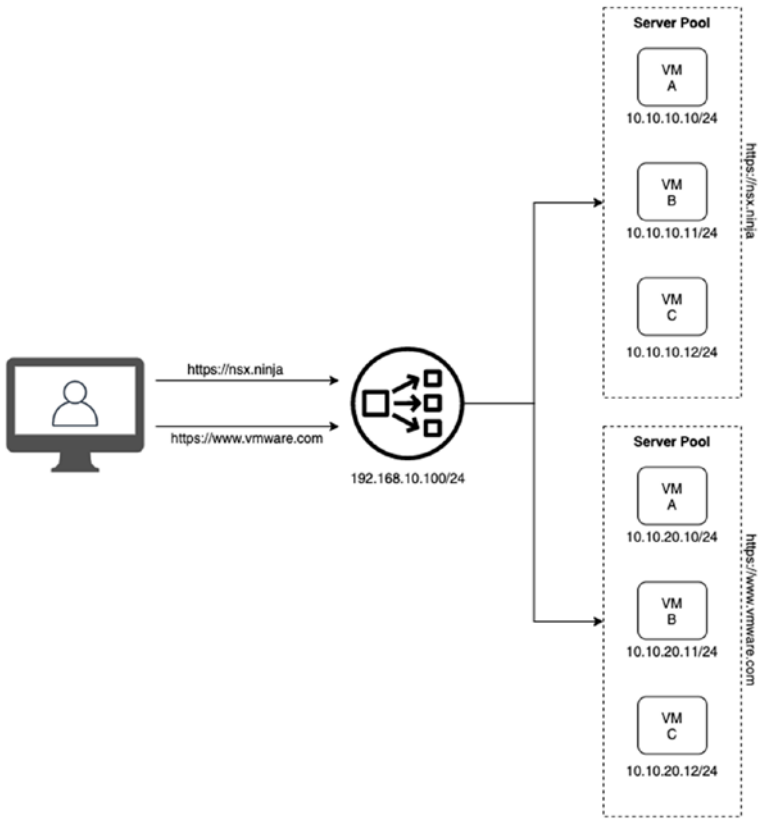


Figure 5-2. Layer-7 load balancing

Load Balancer Component Relation

A load balancer can (or should) be configured on (or attached to) one Tier-1 gateway.

The rules listed in Table 5-1 apply when you deal with a Tier-1 gateway enabled load balancer. See Figure 5-3.

Table 5-1. Rules for a Tier-1 Gateway Enabled Load Balancer

Tier-1 gateway	A load balancer can be attached to only one Tier-1 gateway.
Virtual servers	A load balancer can host one or more virtual servers.
Profiles	One or more virtual servers can use profiles
Server pools	A virtual server can include one or more server pools.
Monitors	One or more server pools can use monitors.

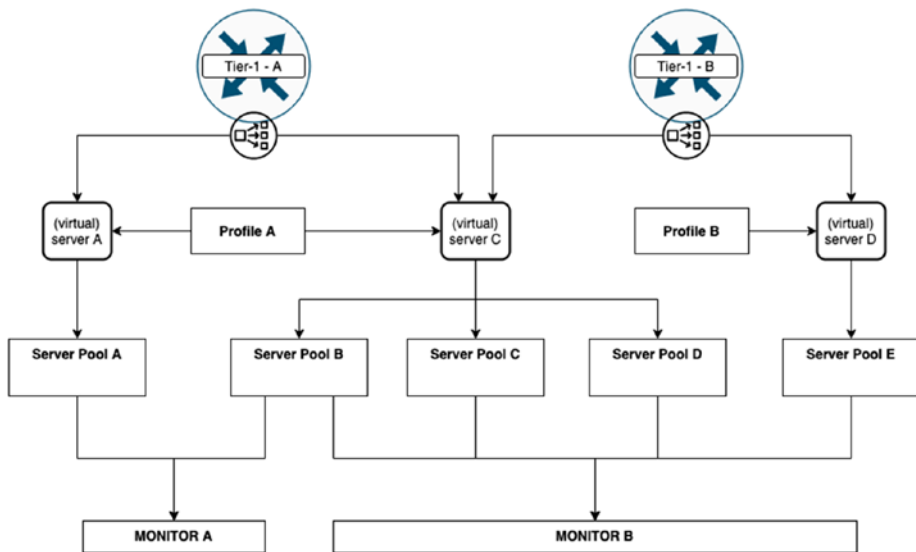


Figure 5-3. Tier-1 gateway enabled load balancer

Load Balancer Architecture

A NSX-T load balancer is attached to (enabled on) a Tier-1 gateway. To enable the load-balancing feature, you need to configure load balancer virtual servers, server pools, load balancer profiles, and load balancer monitors. See Figure 5-4.

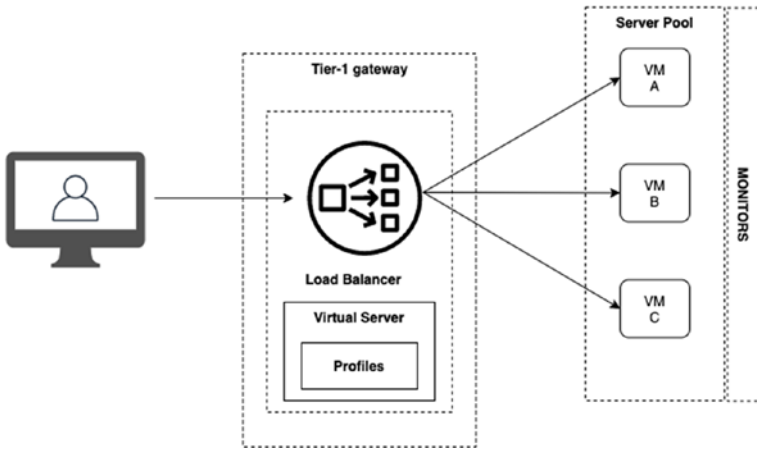


Figure 5-4. Load balancer architecture

Load Balancer Attachment to a Tier-1 Gateway

When a load balancer is enabled on a Tier-1 gateway, a Tier-1 gateway service router (SR) component is instantiated. The Tier-1 gateway SR is always running in Active-Standby HA mode. See Figure 5-5.




Figure 5-5. Tier-1 gateway enabled load balancer

Virtual Servers

A virtual server consists of a virtual IP address, a port, and a protocol. Typically, these are the IP addresses and port numbers the client will perform a connection/session to, and from there, the load will be distributed to the servers in the server pool. See Figure 5-6.

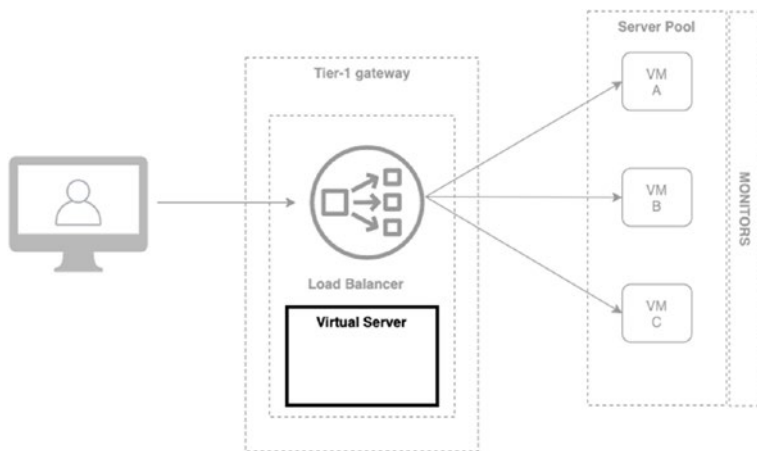


Figure 5-6. *Virtual server*

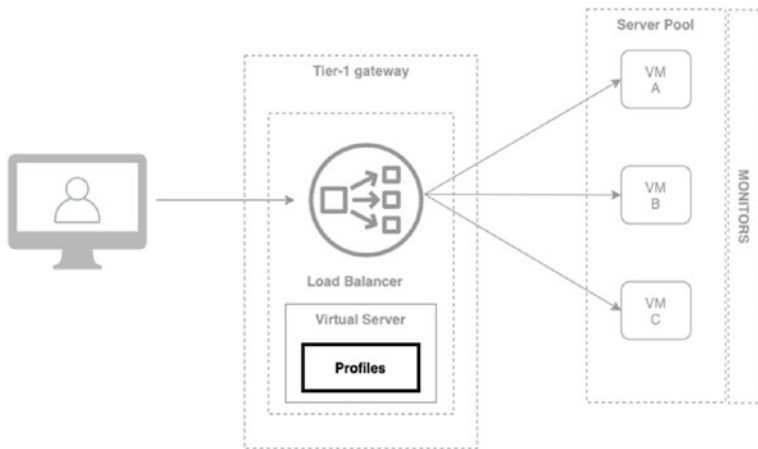
Profiles

A profile is used to configure the behavior and characteristics of the configured virtual servers. Profiles are associated with virtual servers to enhance load-balancing network traffic and simplify traffic-management tasks. See Figure 5-7.

The profiles outlined in Table 5-2 are supported.

Table 5-2. Load Balancer Profiles

Profile Type	Description
Application	This defines how the virtual server processes the network traffic.
Persistence	This is used in stateful applications to redirect all related connections to the same backend server.
SSL	This defines the SSL protocol type and ciphers to be used by the client and server (applicable to Layer-7 load balancers only).

**Figure 5-7.** Load balancer profiles

Server Pools

A server pool contains the servers that offer the service also known as *backend servers*. The servers in the pool typically have the same functionality and services running. The membership of a server pool can be static or dynamic. See Figure 5-8.

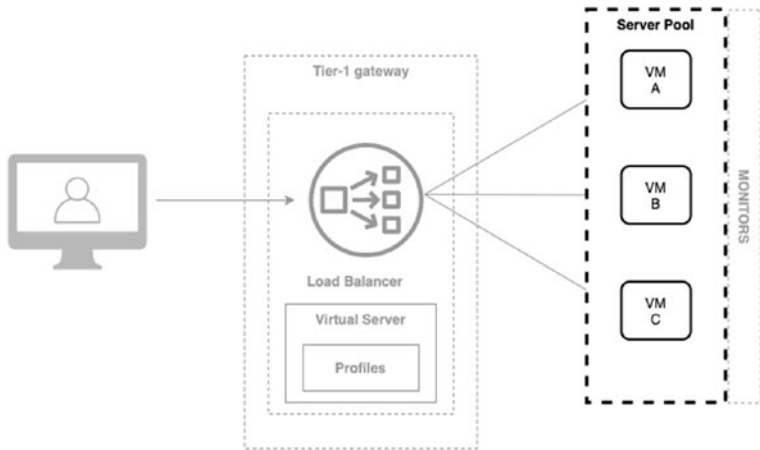


Figure 5-8. Load balancer server pool

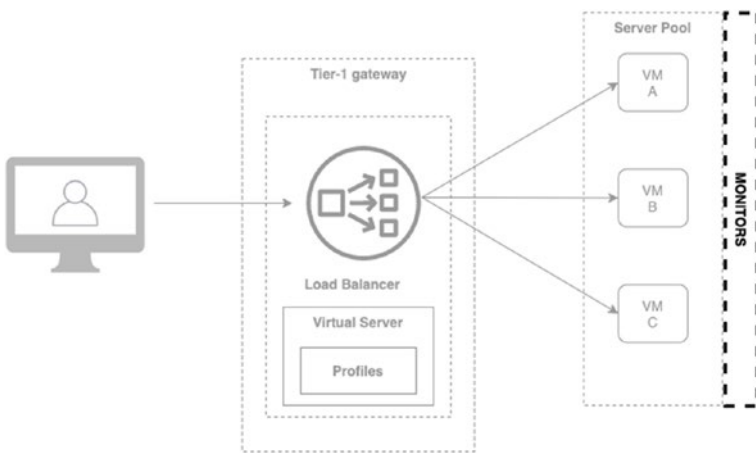
Monitors

A monitor ensures that the servers in the pool are available based on specific values that can measure if an application is really available or not. A session is only forwarded to a member in the server pool when this specific member is available. When the monitor based on its evaluation criteria determines that a specific member in the server pool is not available, no sessions are forwarded to this member until the state changes and the application is available once more on that specific member. See Figure 5-9.

There are two types of load balancer monitors, as described in Table 5-3.

Table 5-3. *Load Balancer Monitors*

Monitor Type	Description
Active	An active monitor is used to test whether a backend server is available.
Passive	A passive monitor is used to check for failures during client connections and will mark servers causing consistent failures as down.

**Figure 5-9.** *Load balancer monitors*

Load Balancer Deployment Modes

Load balancers can be configured in different sizes:

- Small
- Medium
- Large
- Extra large

Depending on the size, the load balancer is capable of hosting different virtual servers and pool members. See Table 5-4.

Table 5-4. *Load Balancer Size Configuration Options*

Quantity Per Load Balancer	Small	Medium	Large	Extra Large
Number of virtual servers per load balancer	20	100	1000	2000
Number of pools per load balancer	60	300	3000	4000
Number of pool members per load balancer	300	2000	7500	10000

The (virtual) edge transport node is available in the following form factors:

- Small
- Medium
- Large
- Extra large

The edge node size determines the number of load balancer instances that it can host. See Table 5-5.

Table 5-5. *Edge Transport Node Size vs Configurable Load Balancers*

Edge Transport Node Size	Small	Medium	Large	Extra
Edge VM - Small	1	Not supported	Not supported	Not supported
Edge VM - Medium	10	1	Not supported	Not supported
Edge VM - Large	40	4	1	Not supported
Edge VM - Extra Large	80	8	2	1
Bare-Metal Edge Node	750	75	18	9

In NSX-T, the load balancer is deployed in two modes: inline mode or one-arm mode.

Inline Mode

In inline mode, the load balancer is *in the traffic path* between the client and the server.

This means that the client and the server *cannot* be connected to the same segment. In this mode, Source NAT (SNAT) is *not* a requirement. See Figure 5-10.




Figure 5-10. Inline mode

One-Arm Mode

In one-arm mode, the load balancer is *not in the traffic path* between the client and the server.

Using one-arm mode, the client and server *can* be anywhere, and thus also hosted on the same segment. In this mode, Source NAT (SNAT) is a requirement. See Figure 5-11.




Figure 5-11. One-arm mode

When you use the one-arm deployment mode, the load balancer will perform Source NAT (SNAT) to make sure that the return traffic coming from the backend servers (from the server pool) to the client is forced through the load balancer (Figure 5-12):

1. The load balancer receives a request from the client (on its configured virtual IP address (VIP)).
2. Depending on the Source NAT (SNAT) configuration, the load balancer replaces the client IP address with the load balancer virtual IP address (VIP).

3. The backend server (from the server pool) sends a response to the load balancer's virtual IP address (VIP).
4. The load balancer forwards the response back to the client.

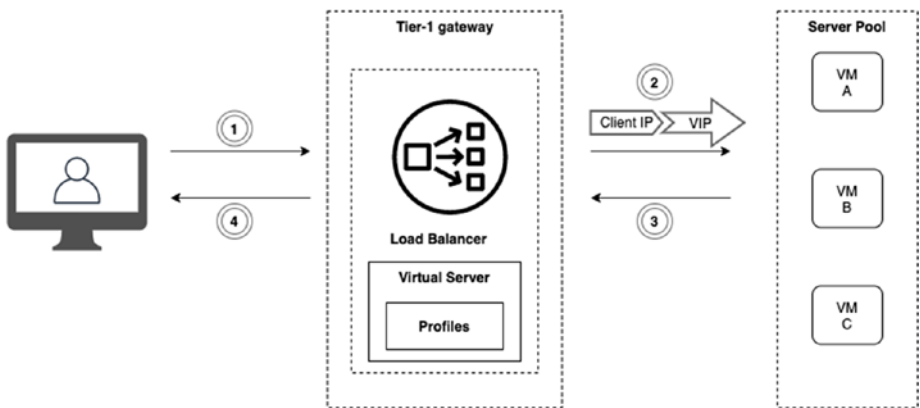


Figure 5-12. One-arm Mode traffic path

Load Balancing Configuration Steps

To configure a load balancer, you need to follow the steps outlined in Table 5-6.

Table 5-6. Load Balancing Configuration Steps

Step	Description
1	Create a Tier-1 gateway.
2	Create a load balancer.
3	Configure the monitor(s).
4	Create the profile(s).
5	Create a server pool.
6	Create a virtual server (VIP).

Load Balancer Creation

To add a load balancer, navigate to Networking ► Load Balancing ► Load Balancers ► Add Load Balancer. When you deploy the load balancer you need to set the parameters outlined in Table 5-7. See Figures 5-13 through 5-16.

Table 5-7. Load Balancer Parameters

Parameter	Description
Name	A friendly name.
Size	The size you want to deploy the load balancer in: Small, Medium, Large, or Extra Large.
Attachment	The Tier-1 gateway that you want to configure the load balancer on (or attach to).

Note When you configure a load balancer on (or attach it to) a Tier-1 gateway, that gateway needs to be attached to a Tier-0 gateway.

The screenshot shows the 'VIRTUAL SERVERS' configuration page in NSX-T. The 'ADD VIRTUAL SERVER' form is displayed with the following details:

- Name:** VIP
- IP Address:** 192.168.10.100 (with a note: e.g. 10.10.10.10 or fc7ef206-db42-)
- Ports:** 80 (with a note: Enter Ports or Port)
- Type:** L4 TCP
- Load Balancer:** LB-T1
- Server Pool:** Selected
- Description:** Enter Description
- Application Profile:** default-tcp-lb-app-profile
- Persistence:** Disabled
- Access List Control:** Configure
- Buttons:** SAVE, CANCEL

Figure 5-13. Virtual server configuration

The screenshot shows the 'LOAD BALANCERS' configuration page in NSX-T. The 'ADD LOAD BALANCER' form is displayed with the following details:

Name	Type	Size	Attachment	Virtual Servers	Status	Alarms
LB-TEST *	Server Load Balancer	Small	T1-Gateway-01	0	Unknown	0

Additional configuration options visible in the form include:

- Description: Enter Description
- Error Log Level: Info
- Tags: Tag, Scope (Max 30 allowed. Click (+) to add.)
- Admin State: Enabled (toggle switch)

NOTE - Before further configurations can be done, fill out mandatory fields above (*), click 'Save' below.

Buttons: SAVE, CANCEL

Figure 5-14. Load balancer configuration (1)

The screenshot shows the 'LOAD BALANCERS' configuration page with a success message:

Load Balancer LB-TEST is successfully created.
Want to continue configuring this Load Balancer?
YES | NO

Figure 5-15. Load balancer configuration (2)

The screenshot shows the 'LOAD BALANCERS' configuration page with a table listing the created load balancer:

Name	Type	Attachment	Virtual Servers	Status	Alarms
LB-TEST	Server Load Balancer	T1-Gateway-01	0	Unknown	0

Additional details for the load balancer:

- Description: Not Set
- Error Log Level: Info
- Tags: 0
- Admin State: Enabled
- VIEW STATISTICS

Figure 5-16. Load balancer configuration (3)

Virtual Server Creation

To add a virtual server, navigate to Networking ► Load Balancing ► Virtual Servers ► Add Virtual Servers. When you create a virtual server, you can select from several protocols, as outlined in Table 5-8. See Figure 5-17.

Table 5-8. *Supported Virtual Server Protocols*

Protocols	Description
L4 TCP	Applications that are running on TCP with Layer-4-only load-balancing requirements
L4 UDP	Applications that are running on UDP with Layer-4-only load-balancing requirements
L7 HTTP	HTTP and HTTPS applications where the load balancer must act based on the Layer-7 parameters





Figure 5-17. *Virtual server protocols*

Layer-4 Virtual Server Creation

When you configure a Layer-4 virtual server, you need to set the parameters outlined in Table 5-9. See Figure 5-18.

Table 5-9. Layer-4 Virtual Server Parameters

Parameter	Description
Name	A friendly name.
IP address	The VIP IP address the client will connect to.
Ports	A supported port (range).
Server pool	The server pool can be created from the wizard or later.
Application profile	<p>The application profile setting is populated by default and is based on the protocol type that you specified when you created the virtual server.</p> <p>Application profiles will define the behavior of a particular network traffic type.</p> <p>The associated virtual server processes network traffic according to the values specified in the application profile.</p>
Persistence profile	<p>The persistence profile is used to ensure the stability of stateful applications.</p> <p>Load balancers implement a persistence so that it directs all related connections to the same server.</p> <p>NSX-T supports different types of persistence to address different types of application needs.</p> <p>Layer-4 virtual servers only support the Source IP option.</p>

**Figure 5-18.** Layer-4 virtual server configuration

Layer-7 Virtual Server Creation

When you configure a Layer-7 virtual server, you need to set the parameters outlined in Table 5-10. See Figures 5-19 and 5-20.

Table 5-10. *Layer-7 Virtual Server Parameters*

Parameter	Description
Name	A friendly name.
IP address	The VIP IP address the client will connect to.
Ports	A supported (single) port Port ranges are not supported when you configure a Layer-7 virtual server.
Server pool	The server pool can be created from the wizard or later.
Application profile	The application profile setting is populated by default and is based on the protocol type that you specified when you created the virtual server. Application profiles will define the behavior of a particular network traffic type. The associated virtual server processes network traffic according to the values specified in the application profile.
Persistence profile	The persistence profile is used to ensure the stability of stateful applications. Load balancers implement a persistence so that it directs all related connections to the same server. NSX-T supports different types of persistence to address different types of application needs. Layer-7 virtual servers support both the Source IP and Cookie persistence options.

(continued)

Table 5-10. (continued)

Parameter	Description
SSL configuration	You can configure SSL parameters on both the server and the client.
Load balancer rules	Layer-7 virtual servers support the configuration of Layer-7 rules.

The screenshot shows the 'VIRTUAL SERVERS' configuration page in the NSX-T interface. The 'ADD VIRTUAL SERVER' form is displayed, showing the following configuration details:

- Name:** VIP
- IP Address:** 192.168.10.100 (with a hint: e.g. 10.10.10.10 or fc7c:f206:db42:)
- Ports:** 80
- Type:** L7 HTTP
- Load Balancer:** LB-TI
- Server Pool:** Sele
- Status:** (indicated by a green dot)
- Alarms:** (indicated by a bell icon)
- Description:** Enter Description
- Application Profile:** default-http-lb-app-profile
- Persistence:** Disabled
- SSL Configuration:** Configure
- Buttons:** SAVE, CANCEL

Figure 5-19. Layer-7 virtual server configuration (1)

The screenshot shows the 'VIRTUAL SERVERS' configuration page in the NSX-T interface. The 'ADD VIRTUAL SERVER' form is displayed, showing the following configuration details:

- Name:** VIP
- IP Address:** 192.168.10.100
- Ports:** 80
- Type:** L7 HTTP
- Load Balancer:** LB-TEST
- Server Pool:** (empty)
- Status:** Success (indicated by a green dot)
- Alarms:** 0 (indicated by a bell icon)
- Description:** Not Set
- Application Profile:** default-http-lb-app-profile
- Persistence:** Disabled
- SSL Configuration:** View
- Buttons:** VIEW STATISTICS

Figure 5-20. Layer-7 virtual server configuration (2)

Application Profile Configuration

To add an application profile, navigate to Networking ► Load Balancing ► Profiles ► Select Profile Type: Application ► Add Application Profile.

Using an application profile, you determine how the virtual servers process network traffic.

By default, there are some default application profiles available that you can choose from, as listed in Table 5-11.

Table 5-11. Load Balancer Default Application Profiles

Default Application Profile	Description
default-http-lb-app-profile	For Layer-7 HTTP virtual servers.
default-tcp-lb-app-profile	For Layer-4 TCP virtual servers. Acceleration is enabled here.
default-udp-lb-app-profile	For Layer-4 UDP virtual servers. Acceleration is enabled here.

You can create user-defined application profiles that are based on TCP, UDP, and HTTP, depending on your requirements. See Figure 5-21.




Figure 5-21. Load balancer default application profiles

Persistence Profile Configuration

To add a persistence profile, navigate to Networking ► Load Balancing ► Profiles ► Select Profile Type: Persistence ► Add Persistence Profile.

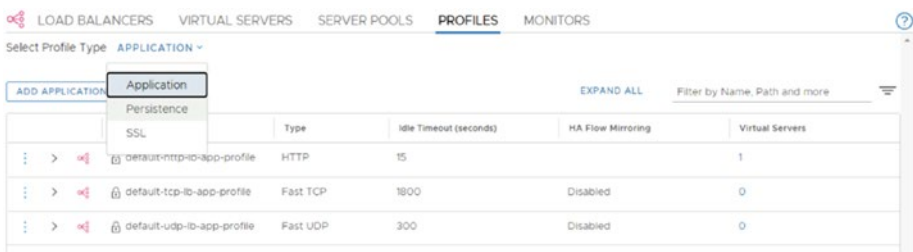
Persistence profiles ensure that session stability is guaranteed when accessing stateful applications. They do this by directing all related connections to the same backend server in the server pool.

The persistence profiles are listed in Table 5-12.

Table 5-12. Load Balancer Persistence Profiles

Profile	Description
Source IP	The source IP tracks sessions based on the client's source IP address. This profile can be used with Layer-4 and Layer-7 virtual servers.
Cookie	The Cookie uses a unique HTTP cookie to identify the session, enabling the client to remain with a server during the session. This profile is used with Layer-7 virtual servers only.

It is possible to create custom persistence profiles that suit your application requirements. A default generic persistence profile is also available. See Figures 5-22 and 5-23.



	Application	Type	Idle Timeout (seconds)	HA Flow Mirroring	Virtual Servers
	default-http-app-profile	HTTP	15		1
	default-tcp-ib-app-profile	Fast TCP	1800	Disabled	0
	default-udp-ib-app-profile	Fast UDP	300	Disabled	0

Figure 5-22. Load balancer profiles

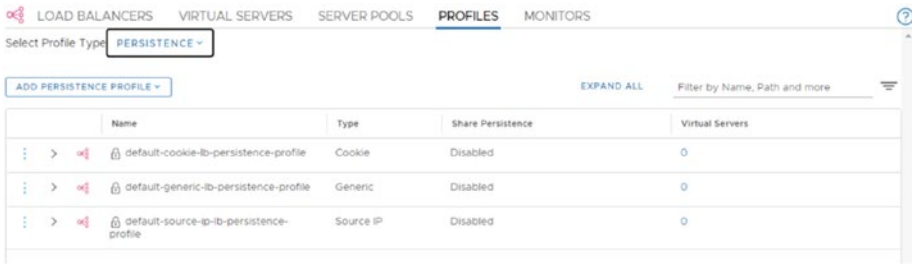


Figure 5-23. Load balancer persistence profile

Layer-7 Load Balancer SSL Modes

The Layer-7 load balancer supports three SSL modes, as outlined in Table 5-13.

Table 5-13. Layer-7 SSL Supported Modes

SSL Mode	Description
Secure Socket Layer (SSL) offload (see Figure 5-24)	The load balancer terminates the client SSL connection (HTTPS). The load balancer uses HTTP to communicate with the servers in the pool.
End-to-end SSL (see Figure 5-25)	The load balancer terminates the client SSL connection (HTTPS). The load balancer creates an SSL connection (HTTPS) toward the servers in the pool.
SSL passthrough	SSL passthrough requires a load balancer rule.

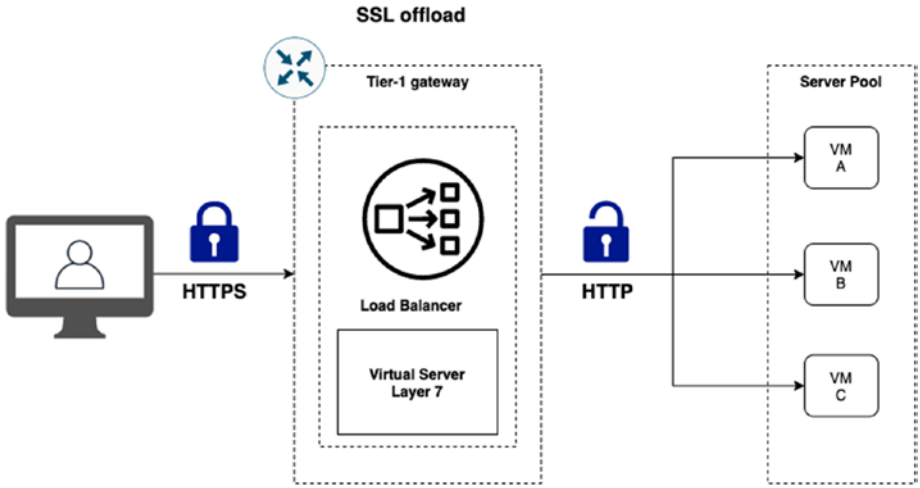


Figure 5-24. SSL offload mode

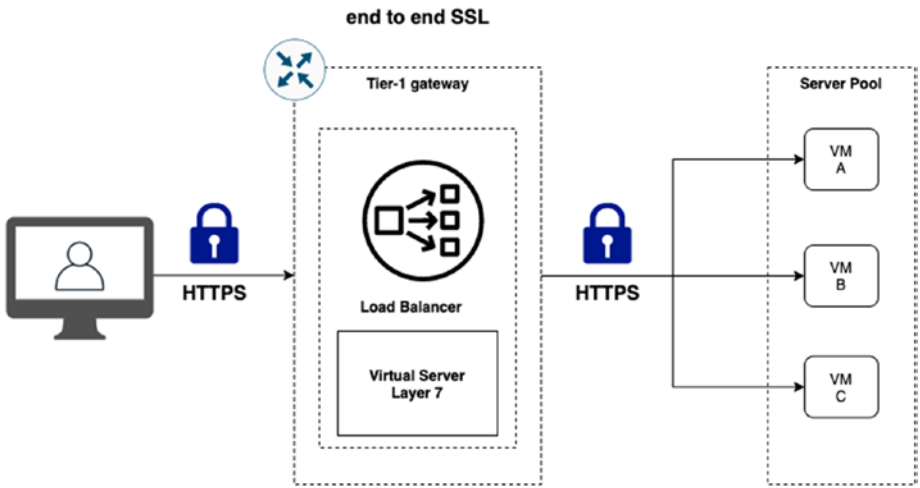


Figure 5-25. End-to-end SSL mode

Layer-7 SSL Profile Configuration

To add an SSL profile, navigate to Networking ► Load Balancing ► Profiles ► Select Profile Type: SSL ► Add SSL Profile.

SSL profiles define the SSL protocol type and ciphers to be used by the client and server.

You can configure two different types of SSL profiles, as outlined in Table 5-14.

Table 5-14. Layer-7 Supported SSL Profiles

SSL Profiles	Description
Client SSL profile (see Figure 5-28)	This is used to configure the SSL connection between the client and the load balancer. This profile is required in SSL offload (see Figure 5-26) and end-to-end SSL (see Figure 5-27) modes.
Server SSL profile (see Figure 5-28)	This is used to configure the SSL connection between the load balancer and the backend server pool. This profile is required in end-to-end SSL mode only.

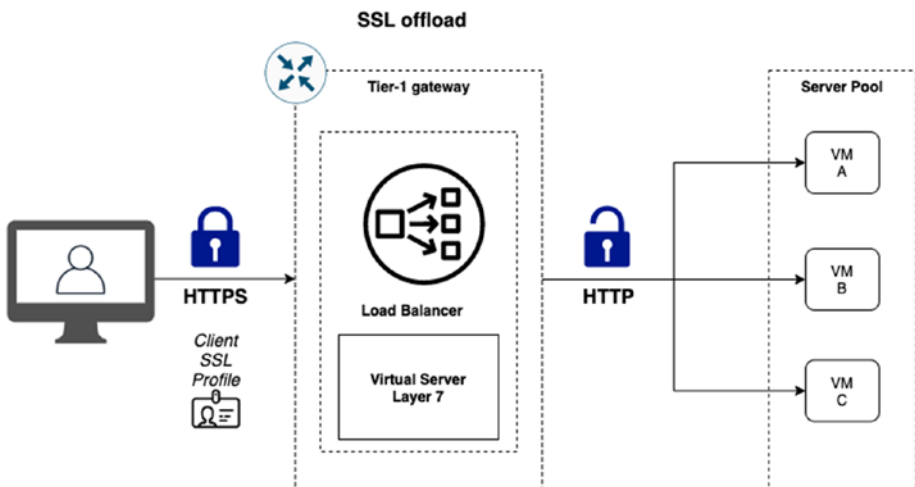


Figure 5-26. SSL offload mode

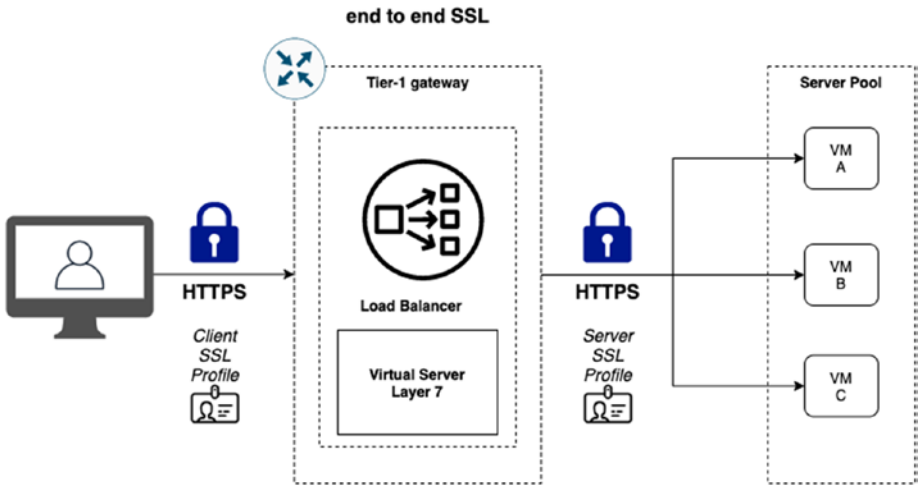


Figure 5-27. End-to-end SSL mode

LOAD BALANCERS VIRTUAL SERVERS SERVER POOLS **PROFILES** MONITORS

Select Profile Type SSL ▾

ADD SSL PROFILE ▾ EXPAND ALL Filter by Name, Path and more

Name	Type	SSL Suite	Session Cache	Virtual Servers
default-balanced-client-ssl-profile	Client SSL Profile	Balanced (recommended)	Enabled	0
default-balanced-server-ssl-profile	Server SSL Profile	Balanced (recommended)	Enabled	0
default-high-compatibility-client-ssl-profile	Client SSL Profile	High Compatibility	Enabled	0
default-high-compatibility-server-ssl-profile	Server SSL Profile	High Compatibility	Enabled	0
default-high-security-client-ssl-profile	Client SSL Profile	High Security	Enabled	0
default-high-security-server-ssl-profile	Server SSL Profile	High Security	Enabled	0

Figure 5-28. SSL profiles

Layer-7 SSL Load Balancer Rules Configuration

To add a load-balancing rule, navigate to Networking ► Load Balancing ► Virtual Servers ► Select a Virtual Server ► Go to the Load Balancer Rules Section.

When you configure a Layer-7 virtual server, you can optionally configure additional load balancer rules to customize the load-balancing behavior.

You have the option to configure Match and Action rules. Load balancer rules support RegEx to configure the match conditions. See Figures 5-29 and 5-30.

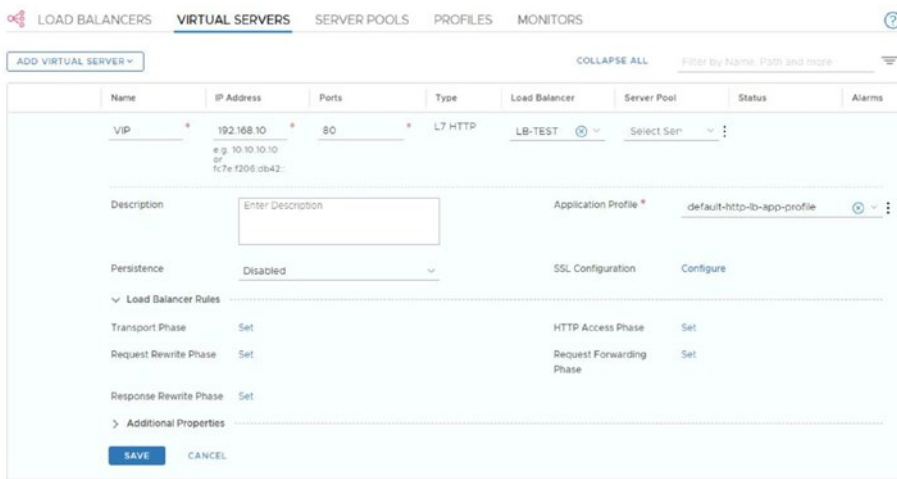


Figure 5-29. Layer-7 virtual server configuration - load balancer rules




Figure 5-30. Load balancer rules

Server Pool Creation

To add a server pool, navigate to Networking ► Load Balancing ► Server Pools ► Add Server Pool.

When you configure a server pool, you need to set the parameters outlined in Table 5-15. See Figures 5-31 through 5-33.

Table 5-15. Load Balancer Server Pool Parameters

Parameter	Description
Name	A friendly name.
Load-balancing algorithm	The load-balancing algorithms control the way that the incoming connections are distributed across the servers in the server pool.
Pool members/group	These members can be static or dynamic if configured as a group.
SNAT translation mode	The load balancer receives the traffic from the server that is destined to the client. SNAT can be enabled per server pool.
Active monitor	The active monitors support verification checks based on HTTP, HTTPS, ICMP, TCP, and UDP.

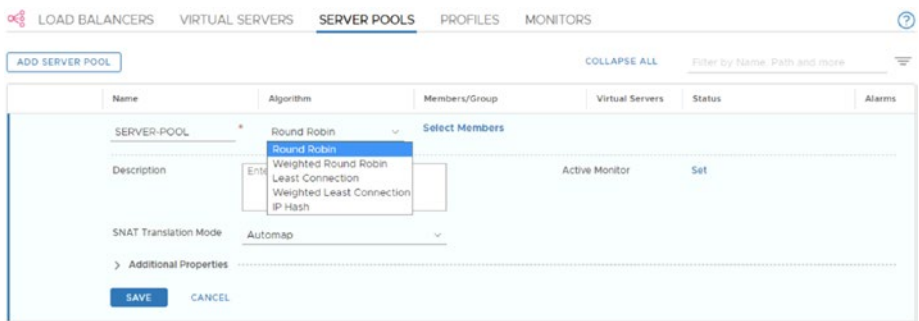


Figure 5-31. Load balancer server pool configuration (1)

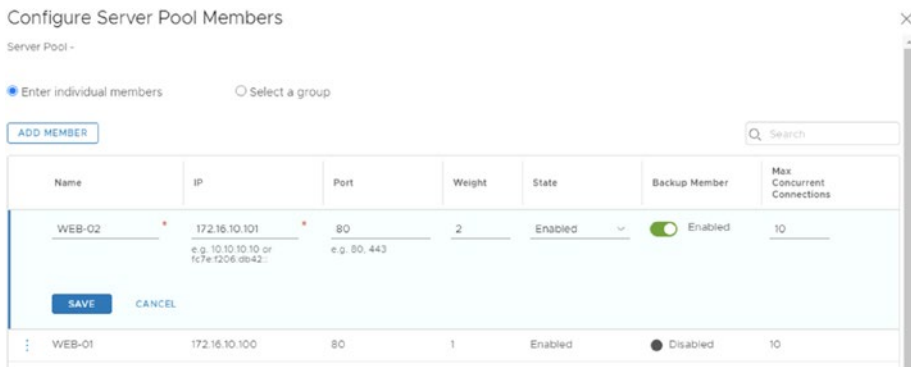


Figure 5-32. Load balancer server pool configuration (2)




Figure 5-33. Load balancer server pool configuration (3)

Load-Balancing Algorithms

Load-balancing algorithms control the way the incoming connections are distributed across the servers in the server pool.

Table 5-16 explains the load-balancing algorithms.

Table 5-16. *Load-Balancing Algorithms*

Load-Balancing Methods
Algorithms

Round robin	Client requests are cycled through the available servers.
Weighted round robin	Each server is assigned a weight based on its performance. This weight determines how many client requests the server receives compared to other servers in the pool.
Least connection	New connections are sent to the server that has the fewest connections.
Weighted least connection	Each server is assigned a weight based on its connection capacity. This weight determines how many client requests the server receives compared to other servers in the pool.
IP hash	A server is selected on a hash of the source IP address and the total weight of all running servers.

SNAT Translation Mode Configuration

Depending on your load-balancing topology, Source NAT (SNAT) might be required so that the load balancer can receive the traffic from the server that is destined to the client. SNAT can be enabled per server pool. See Figure 5-34.

When creating a server pool, the SNAT translation modes outlined in Table 5-17 are available.

Table 5-17. Supported SNAT Translation Modes

SNAT Translation Mode	Description
Automap	The load balancer uses its interface IP address and ephemeral port to continue the communication with a client who initially connected to one of the server’s established listening ports.
Disabled	No SNAT.
IP pool	Allows users to specify a single IP address range to be used by servers for SNAT.

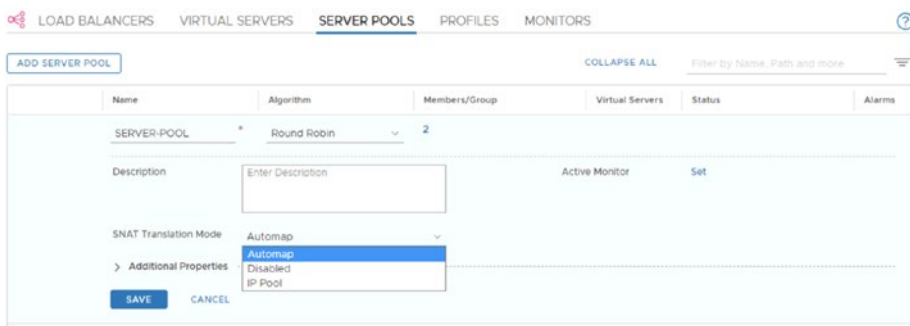


Figure 5-34. SNAT translation modes

Active Monitoring Configuration

You can use active monitors to verify whether the backend servers (in the server pool) are available.

The active monitors support verification checks based on HTTP, HTTPS, ICMP, TCP, and UDP.

To add an active monitor, navigate to **Networking** ➤ **Load Balancing** ➤ **Monitors** ➤ **Select Monitor Type: Active** ➤ **Add Active Monitor**. See [Figure 5-35](#).

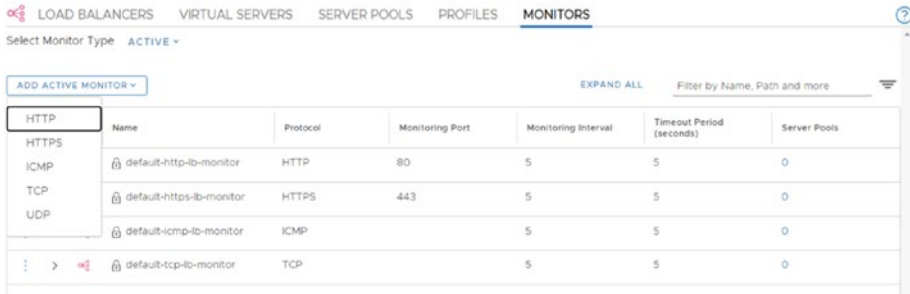


Figure 5-35. Active monitor configuration

Passive Monitoring Configuration

You can also use passive monitors to check for errors during client connections and mark servers causing consistent failures as “down”.

The number of failures is incremented when the conditions listed in Table 5-18 occur.

Table 5-18. Passive Monitor Options

Option	Description
Layer-7 virtual servers	Failures related to TCP connection errors or SSL handshakes
Layer-4 TCP virtual servers	Failures related to TCP connection errors
Layer-4 UDP virtual servers	Internet Control Message Protocol (ICMP) errors

To add a passive monitor, navigate to Networking ► Load Balancing ► Monitors ► Select Monitor Type: Passive ► Add Passive Monitor. See Figures 5-36 and 5-37.




Figure 5-36. Passive monitor configuration (1)

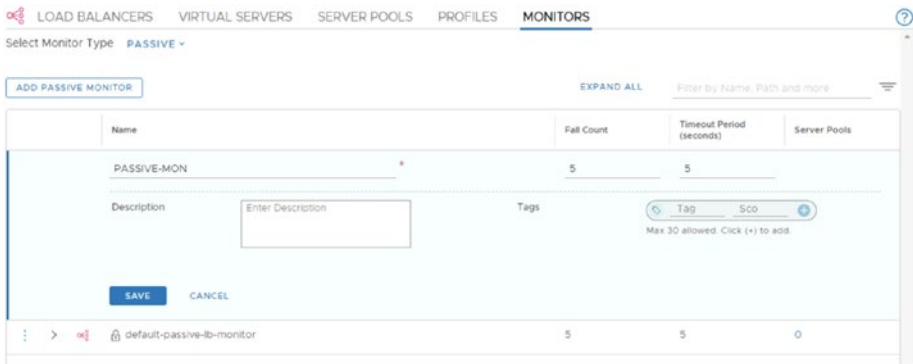


Figure 5-37. Passive monitor configuration (2)

Summary

In this chapter, you learned about the load balancer architecture and components. You should now be able to identify the load balancer deployment modes and create and configure a load balancer using Layer-4 and Layer-7 virtual servers.

In the next chapter, you’ll learn about the different VPN types that NSX-T offers.

CHAPTER 6

NSX-T VPN

This chapter describes the requirements for IPsec VPN using NSX-T and identifies the different VPN types. It also describes how to create and configure Layer-3 IPsec VPN tunnels and Layer-2 VPN (L2VPN) tunnels.

NSX-T VPN Services

NSX-T supports Layer-3 IPsec VPN (route-based and policy-based) and Layer-2 services.

IPsec and Layer-2 VPN types are supported on Tier-1 and Tier-0 gateways and require NSX Edge transport nodes to operate.

IPsec VPN Use Cases

The use cases for IPsec VPN are listed here:

- IPsec VPN allows you to connect two remote IP networks together in a secure manner.
- IPsec extends IP network connectivity between two or more sites.
- Layer-3 IPsec VPN is route-based and allows connectivity between different, non-overlapping IP subnets.
- IPsec provides a secure communication channel when used over protocols, like Generic Routing Encapsulation (GRE).

In Figure 6-1, you can see the following:

1. An IPsec connection between two sites where NSX-T is not used.
2. An IPsec connection between two sites where NSX-T is used on only one site.
3. An IPsec connection between two sites where NSX-T is used on both sites.

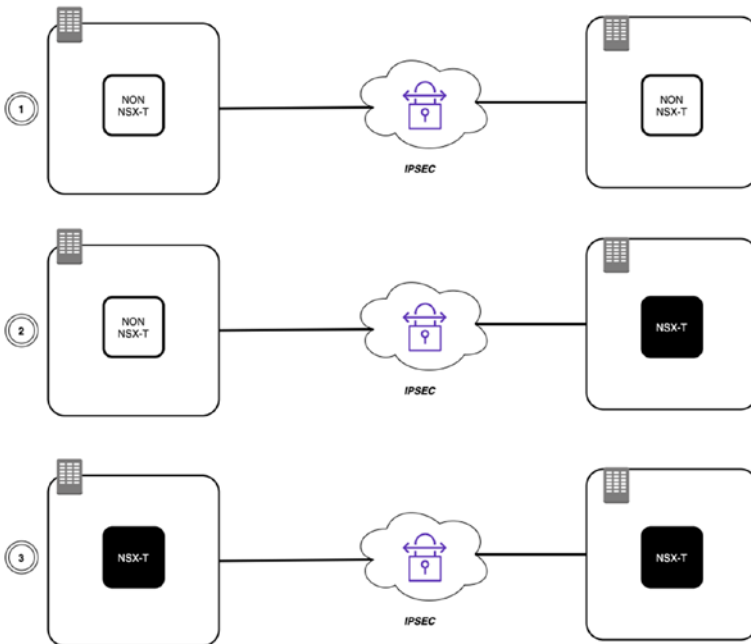


Figure 6-1. IPSEC VPN use cases

IPsec VPN Methods

IPsec by itself is not a single standalone protocol. It is a collection of protocols designed to provide confidentiality, authentication, and integrity to a VPN connection.

To make this happen, IPsec uses Internet Key Exchange (IKE). IKE manages the connection to a peer, defines the security associations used to secure and validate data exchanges, and defines security protocols used to carry IP traffic over the VPN.

Security Associations (SA) are the basic component of IPsec; they contain information about the security parameters that are negotiated/exchanged between two peers. The two SAs that are available are IKE (or ISAKMP) SA and IPsec SA.

The IKE SA is used in the control plane of the VPN and contains a combination of mandatory and optional values, as listed in Table 6-1.

Table 6-1. *IKE SA Values*

Optional Value	Mandatory/Optional
Encryption algorithm	Mandatory
Hash algorithm	Mandatory
Authentication method	Mandatory
Diffie-Hellman group	Mandatory
Lifetime	Optional

IPsec VPN tunnel packets can use the following headers:

- Authentication header (AH)
- Encapsulating security payload header (ESP)

The authentication header is used for authentication, whereas the encapsulating security payload header enables the encryption of the protected payload.

IPsec VPN Modes

Two VPN modes are supported in IPsec, as defined in Table 6-2. See Figure 6-2.

Table 6-2. *IPsec VPN Modes*

Mode	Description
Transport mode	Preserves the original IP header and is typically used for remote-access VPNs.
Tunnel mode	Encapsulates the entire IP datagram with a new header and is typically used for site-to-site VPNs between IPsec-enabled gateways.

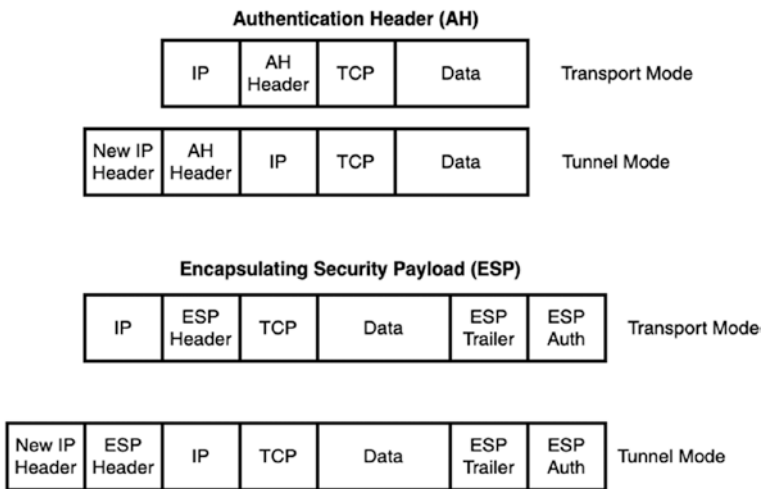


Figure 6-2. *AH and ESP headers*

IPsec Protocols and Algorithms

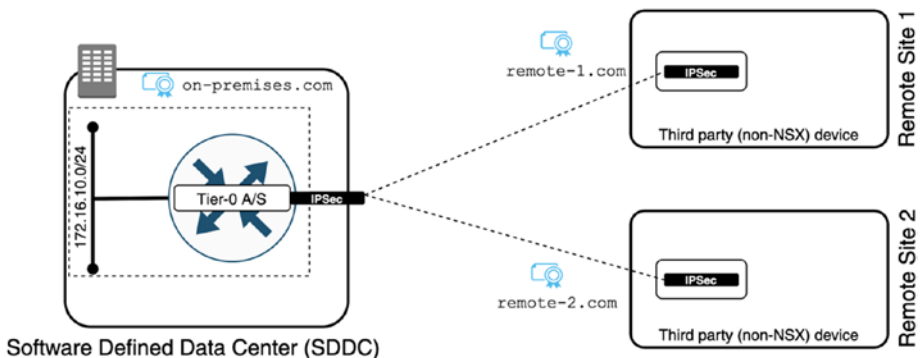
IPsec VPN includes several protocols and algorithms, as explained in Table 6-3.

Table 6-3. VPN Protocols and Algorithms

Protocol/Algorithm	Description
Key Management	IPsec VPN uses the Internet Key Exchange (IKE) protocol to negotiate security parameters. IKE runs by default over UDP/500 port. If NAT is detected in the gateway, the port is set to UDP/4500. IKEv1 (RFC 2409) and IKEv2 (RFC 5996) are both supported.
Authentication	This can be done using pre-shared keys or certificates.
Encryption	The supported encryption algorithm in NSX-T is Advanced Encryption Standard (AES).
Data integrity	The supported data integrity algorithm in NSX-T is Secure Hash Algorithm (SHA).

IPsec VPN Certificate-Based Authentication

With IPsec VPN certificate-based authentication, there is a certificate in place for each local endpoint. For each local endpoint, there is a trusted certificate authority (CA) including its full certificate chains. Pre-shared keys (PSKs) are used to improve manageability. A certificate revocation list (CRL) is also used. See Figure 6-3.

**Figure 6-3.** IPsec VPN certificate-based authentication

IPsec VPN Dead Peer-to-Peer Detection

dead peer detection (DPD) is a mechanism that detects whether an IPsec connection is alive or dead. An NSX-T DPD profile specifies the number of seconds to wait between the probes to detect if an IPsec peer is alive or dead.

In NSX-T, there is a default system-generated DPD profile in place, called `nsx-default-l3vpn-dpd-profile`. This profile is assigned when you configure an IPsec VPN service. If you do not want to use the default profile, you can use the NSX UI to create your own custom DPD profile.

You can look at the default DPD profile or create a custom one by navigating to **Networking** > **Network Services** > **VPN** > **Profiles** > **Select Profile Type** > **DPD Profiles**. See Figure 6-4.

The screenshot shows the NSX-T Profiles page. The breadcrumb navigation is: **VPN SERVICES** > **IPSEC SESSIONS** > **L2 VPN SESSIONS** > **LOCAL ENDPOINTS** > **PROFILES**. Below the navigation, there is a dropdown menu for "Select Profile type:" with "DPD PROFILES" selected. There is an "ADD DPD PROFILE" button and an "EXPAND ALL" button. A filter option "Filter by Name, Path and more" is also visible. The main content is a table with the following data:

Name	DPD Probe Mode	DPD Probe Interval (seconds)	Retry Count	Sessions	Status
<code>nsx-default-l3vpn-dpd-profile</code>	Periodic	60	10	0	Success

Figure 6-4. Dead peer detection (DPD)

IPsec VPN Types

The IPsec types that are supported by NSX-T are policy-based and route-based, as described in Table 6-4.

Table 6-4. *IPsec VPN Types*

IPsec VPN Type	Characteristics
Policy-based VPN	<p>A VPN policy determines the IPsec-protected traffic going through the VPN tunnel.</p> <p>You need a static configuration modification of the VPN policy when the network topology is changing.</p>
Route-based VPN	<p>Protected local and remote networks are learned dynamically using BGP routing exchanges. Endpoint routes are learned through a specific interface called a Virtual Tunnel Interface (VTI). All packets routed through the VTI are protected by IPsec.</p> <p>Tunnel redundancy is supported because:</p> <ul style="list-style-type: none"> – Remote subnets can be learned over multiple BGP peers. – The primary tunnel is active, whereas the secondary tunnel is standby. – If the peer is unreachable on the primary tunnel, the secondary tunnel becomes active. <p>OSPF is not supported.</p>

IPsec VPN Deployment Considerations

When you deploy IPsec VPN, you should consider the following:

- Layer-3 IPsec VPN services are available on Tier-1 and Tier-0 gateways.
- Protected networks must be segments created through the NSX-T UI or policy APIs.
- Segments can be connected to either Tier-0 or Tier-1 gateways to use the configured VPN services.

- When you are using tenants with overlapping networks, this will require NAT on the Tier-0 gateways.
- The NSX-T data center supports site-to-site IPsec VPNs in tunnel mode.
- IPsec tunnels use DPDK to accelerate performance.

IPsec High Availability (HA)

IPsec VPN supports being highly available in an Active/Standby mode on Tier-1 and Tier-0 gateways. The VPN services use a gateway-level failover and the VPN configuration data is synchronized. The VPN “state” is not synchronized because tunnels are reestablished on failover. You can use high-availability virtual IP addresses for external connections.

Figure 6-5 shows a simulation where the IPsec tunnel is active on the Tier-0 gateway (1) on the edge transport node 1. When the edge transport node 1 becomes unavailable (2) for whatever reason, the Standby edge transport node 2 will reestablish the connection (3).

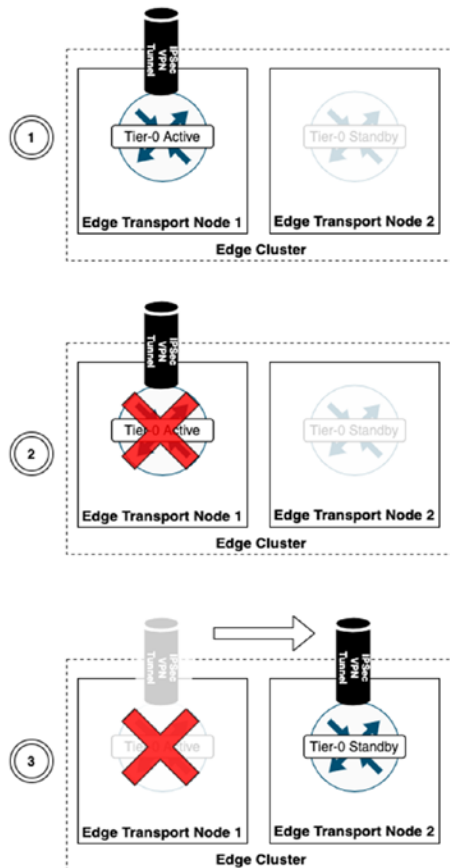


Figure 6-5. IPsec HA

IPsec VPN Scalability

Table 6-5 lists the different sizes of NSX-T Edge transport nodes and the number of VPN sessions that are supported per size.

The numbers are soft limits or tested maximums but may change depending on the NSX-T version. More on the maximums can be found at <https://configmax.vmware.com/>.

Table 6-5. *IPsec VPN Scalability*

	Small	Medium	Large	Extra Large	Bare Metal
Number of tunnels per session (policy-based)	N/A	256	256	256	256
Number of sessions	128	256	512	512	512
Number of IPsec tunnels	1024	2048	4096	4096	4096

IPsec VPN Configuration

When you want to configure a Layer-3 IPsec VPN tunnel, you need to perform the configuration steps in Table 6-6.

Table 6-6. *Layer-3 IPsec VPN Tunnel Configuration Steps*

Step	Description
1	Create an IPsec VPN service (and apply this on an existing gateway).
2	(Optional) Configure a DPD profile.
3	(Optional) Configure an IKE profile.
4	(Optional) Configure an IPsec profile.
5	Add a local endpoint.
6	Configure an IPsec session: <ul style="list-style-type: none"> – Configure route-based IPsec. – Configure policy-based IPsec.

IPsec VPN Service Configuration

You can create an IPsec VPN service by navigating to Networking ► Network Services ► VPN ► VPN Services ► Add Service: IPsec. See Figures 6-6 through 6-8.



Figure 6-6. IPsec VPN service configuration (1)

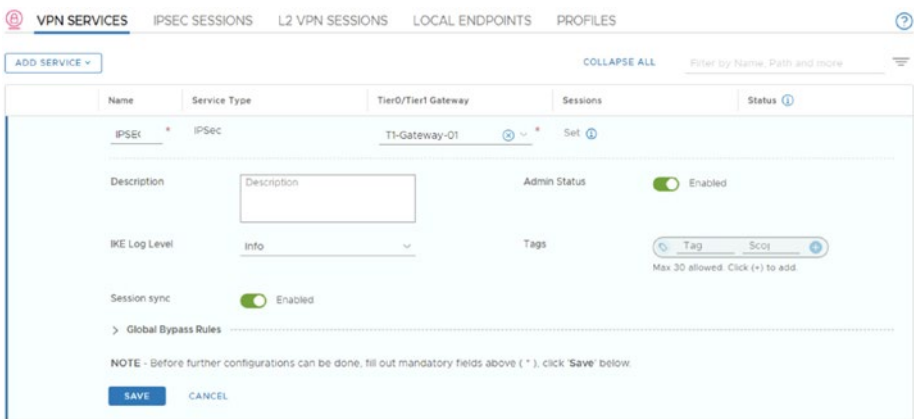


Figure 6-7. IPsec VPN service configuration (2)




Figure 6-8. IPsec VPN service configuration (3)

When you configure a new IPsec service, you can specify the parameters outlined in Table 6-7.

Table 6-7. *IPsec Service Configuration Parameters*

Parameter	Description
Name	A name for the IPsec service.
Tier-0/Tier1 gateway	Select a Tier-0/Tier-1 gateway to associate this IPsec VPN service to.
IKE log level	Enable VPN service logging. The IKE logging level determines the amount of information that you want collected for the IPsec VPN traffic. The default is set to the Info level.
Admin status	Enable or disable the IPsec VPN service. By default, the value is set to Enabled to enable the service on the Tier-0 gateway.
Tags	A user-defined tag for references.

DPD Profile Configuration

You can create an DPD profile by navigating to Networking ► Network Services ► VPN ► Profiles ► Select Profile Type: DPD Profile. See Figure 6-9.

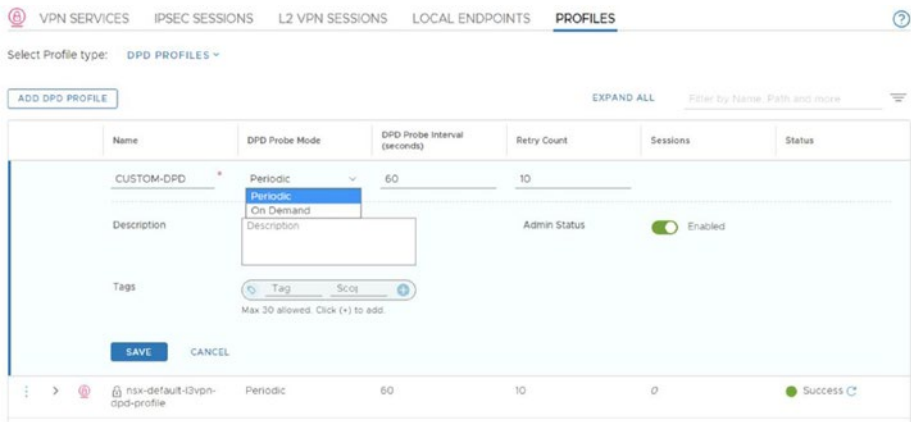


Figure 6-9. DPD profile configuration

When you configure a new DPD profile, you can specify the parameters outlined in Table 6-8.

Table 6-8. DPD Profile Configuration Parameters

Parameter	Description
Name	A name to identify the service.
DPD probe mode	Periodic: For a periodic DPD probe mode, a DPD probe is sent every time the specified DPD probe interval time is reached. On Demand: For an on-demand DPD probe mode, a DPD probe is sent if no IPsec packet is received from the peer site after an idle period. The value in the DPD Probe Interval text box determines the idle period used.
DPD probe interval (sec):	A value in seconds that defines at which times a DPD detection packet should be sent.
Retry count	The number of allowed retries.
Admin status	Enable or disable the profile.
Tags	A user-defined tag for references.

IKE Profile Configuration

You can create an IKE profile by navigating to Networking ► Network Services ► VPN ► Profiles ► Select Profile Type: IKE Profile. See Figures 6-10 and 6-11.

The screenshot shows the 'PROFILES' tab in the NSX-T interface. Under 'Select Profile type: IKE PROFILES', there is a table listing several IKE profiles. Each row includes a name, IKE version, encryption and digest algorithms, Diffie-Hellman group, session count, and status.

Name	IKE Version	Encryption Algorithm	Digest Algorithm	Diffie-Hellman	Sessions	Status
CNSA	IKE V2	AES 256	SHA2 384	Group 20	0	Success
FIPS	IKE FLEX	AES 128	SHA2 256	Group 20	0	Success
Foundation	IKE V1	AES 128	SHA2 256	Group 14	0	Success
nsx-default-2vpn-ik...	IKE V2	AES 128	SHA2 256	Group 14	0	Success
nsx-default-3vpn-ik...	IKE V2	AES 128	SHA2 256	Group 14	0	Success
PRIME	IKE V2	AES GCM 128	Not Set	Group 19	0	Success
Suite-B-GCM-128	IKE V2	AES 128	SHA2 256	Group 19	0	Success
Suite-B-GCM-256	IKE V2	AES 256	SHA2 384	Group 20	0	Success

Figure 6-10. IKE profile configuration (1)

The screenshot shows the configuration form for a new IKE profile. The form includes fields for Name, IKE Version, Encryption Algorithm, Digest Algorithm, Diffie-Hellman, Sessions, and Status. The 'Name' field is filled with 'IKE', 'IKE Version' is 'IKE V2', 'Encryption Algorithm' is 'AES 128', 'Digest Algorithm' is 'SHA2 256', and 'Diffie-Hellman' is 'Group 14'. The 'SA Lifetime (seconds)' is set to '86400'. There are also fields for Description and Tags.

Figure 6-11. IKE profile configuration (2)

When you configure a new IKE profile, you can specify the parameters outlined in Table 6-9. See Figure 6-12.

Table 6-9. *IKE Profile Configuration Parameters*

Parameter	Description
Name	A name to identify the service.
IKE version	Select IKE V1, IKE V2, or IKE FLEX, depending on your technical requirements.
Encryption algorithm	The encryption algorithm used during the Internet Key Exchange (IKE) negotiation.
Digest algorithm	The secure hashing algorithm used during the IKE negotiation.
Diffie-Hellman	The cryptography schemes that the peer site and the NSX Edge instance use to establish a shared secret over an insecure communications channel.
SA lifetime (sec)	The lifetime (in seconds) of the security associations (individual communicating peer identifiers), after which a renewal is required.
Tags	A user-defined tag for references.

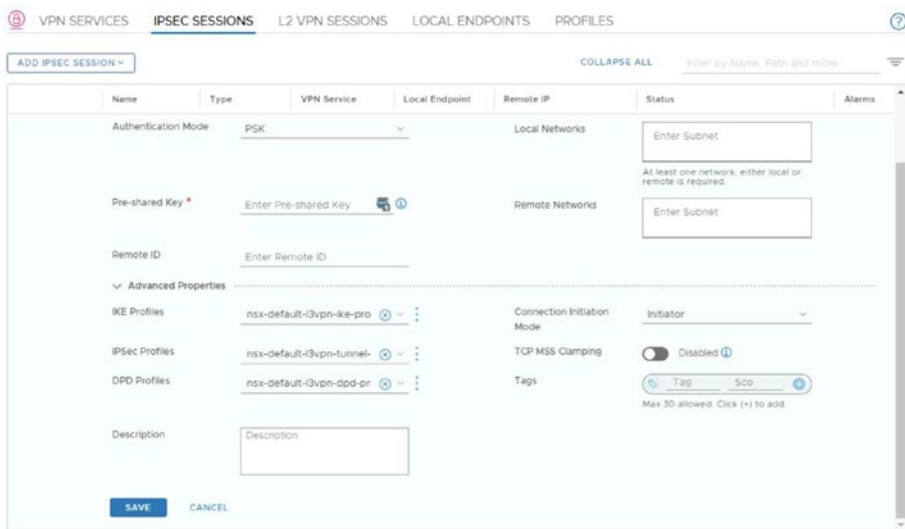


Figure 6-12. IPsec Session Configuration

IPsec Profile Configuration

You can create an IPsec profile by navigating to Networking ► Network Services ► VPN ► Profiles ► Select Profile Type: IPsec Profile. See Figures 6-13 and 6-14.

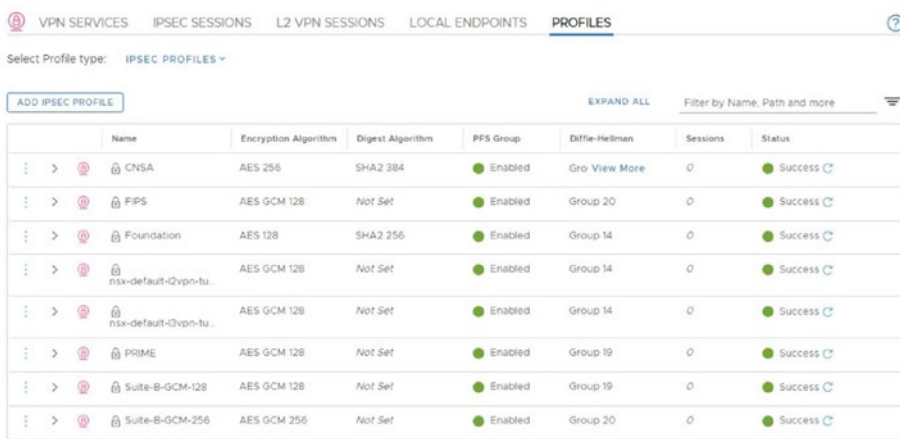


Figure 6-13. IPsec profiles (1)

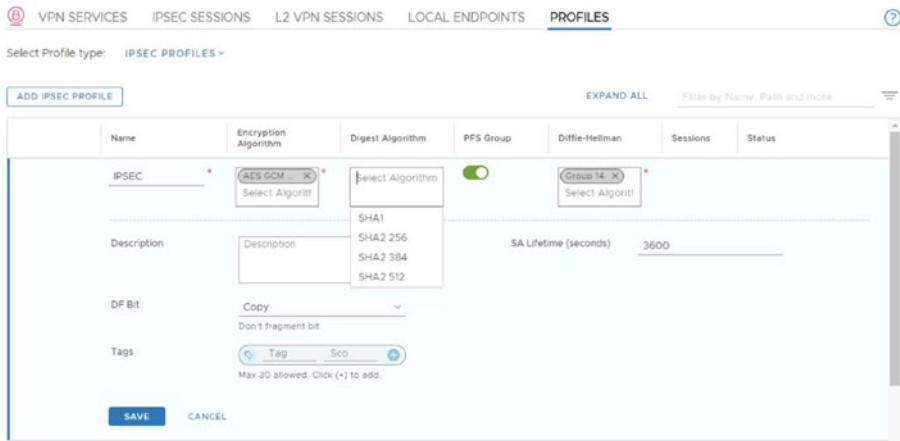


Figure 6-14. IPsec profiles (2)

When you configure a new IPsec profile, you can specify the parameters outlined in Table 6-10.

Table 6-10. IPsec Profile Configuration Parameters

Parameter	Description
Name	A name to identify the service.
Encryption algorithm	The encryption algorithm used during the IPsec negotiation.
Digest algorithm	The secure hashing algorithm used during the IPsec negotiation.
PFS group	The Perfect Forward Secrecy (PFS) group, which adds protection to the keys that build secure channels. You can enable or disable this option.
Diffie-Hellman	The cryptography schemes that the peer site and NSX Edge use to establish a shared secret over an insecure communications channel.

(continued)


Table 6-10. (continued)

Parameter	Description
SA lifetime (sec)	The lifetime (in seconds) of the security associations (individual communicating peer identifiers), after which a renewal is required.
DF bit	Determines if the encrypted traffic should copy the DF (don't fragment) bit from the inner payload to the encrypted traffic.
Tags	A user-defined tag for references.

IPsec VPN Session Configuration

When you configure an IPsec session, you can choose between two VPN types—policy-based or route-based.

You can create an IPsec VPN session by navigating to Networking ► Network Services ► VPN ► IPsec Sessions ► Add IPsec Sessions. See Figure 6-15.

**Figure 6-15.** IPsec sessions

Policy-Based IPsec Session

When you select the policy-based option, IPsec tunnels are used to connect multiple local subnets that are behind the NSX Edge transport node, with peer subnets on the remote VPN site. See Figure 6-16.

The screenshot shows the configuration page for a new IPsec session in NSX-T. The navigation tabs include VPN SERVICES, IPSEC SESSIONS (selected), L2 VPN SESSIONS, LOCAL ENDPOINTS, and PROFILES. The session name is 'POLICY' and the type is 'Policy Based'. The VPN Service is set to 'Select', Local Endpoint to 'Se', and Remote IP to 'Enter Remote IP'. The Compliance suite is 'None' and the Authentication Mode is 'PSK'. The Pre-shared Key field is empty with a help icon. The Remote ID field is 'Enter Remote ID'. The Local Networks and Remote Networks fields are empty with a note: 'At least one network, either local or remote is required.' The Status is 'Enabled' with a green toggle. There are 'SAVE' and 'CANCEL' buttons at the bottom.

Figure 6-16. Policy-based IPsec session

When you configure a new VPN session (policy), you can specify the parameters outlined in Table 6-11.

Table 6-11. Policy-Based IPsec Session Configuration Parameters

Parameter	Description
Name	A name to identify the service.
Type	This is defined in the previous selection.
VPN service	This is the predefined service to use with this session.
Local endpoint	This is the earlier configured local endpoint for use with this configuration session.
Remote IP	The specifies the IP address of the remote IPsec-capable gateway for building the secure connection.
Authentication mode	This defines whether to use the pre-shared key (PSK) or certificate-based connection authentication.
Local networks and remote networks	This defines the interesting traffic that should be encrypted through this VPN session.

(continued)

Table 6-11. (continued)

Parameter	Description
Pre-shared key	This specifies the string to define the key if the authentication mode is PSK.
Remote ID	This defines the identifier of the remote peer for verifying peering authenticity.
Tags	A user-defined tag for references.

In the Advanced Properties area, you can select the predefined DPD, IKE, and IPsec profiles from the drop-down menus. You can also define which side initiates the connection.

Route-Based IPsec Session

When you select the route-based option, tunneling is provided on traffic that is based on routes. These routes were learned dynamically over a virtual tunnel interface (VTI) by using a preferred protocol such as BGP. IPsec secures all the traffic flowing through the VTI. See Figure 6-17.

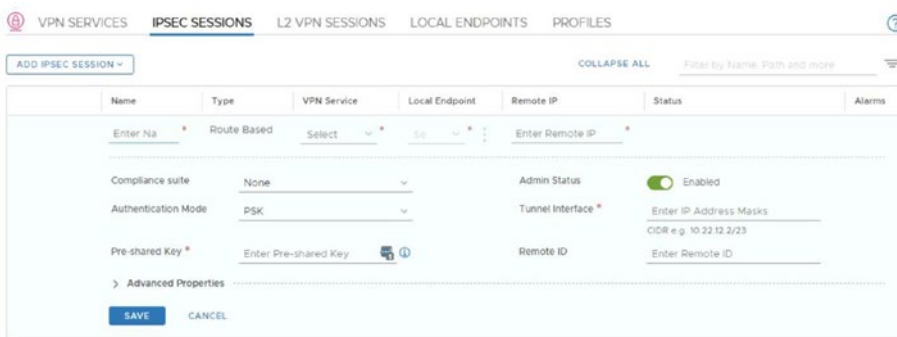


Figure 6-17. Route-based IPsec session

When you configure a new VPN session (route), you can specify the parameters outlined in Table 6-12.

Table 6-12. *Route-Based IPsec Session Configuration Parameters*

Parameter	Description
Name	A name to identify the service.
Type	This is defined in the previous selection.
VPN service	This is the predefined service to use with this session.
Local endpoint	This defines an earlier configured local endpoint to use with this session configuration.
Remote IP	This defines the IP address of the remote IPsec-capable gateway for building the secure connection.
Compliance suite	You can specify a security compliance suite such as CNSA, FIPS, Foundation, PRIME, or Suite-B to configure the security profiles used for an IPsec VPN session.
Authentication mode	This defines whether to use a pre-shared key (PSK) or certificate-based connection authentication.
Admin status	This enables and disables the service.
Tunnel interface	This defines the IP address of the local virtual tunnel interface (VTI) that is created to use with this session.
Pre-shared key	This provides the string for defining the key if the authentication mode is PSK.
Remote ID	This specifies the identifier of the remote peer for verifying the authenticity of the peering.
Tags	A user-defined tag for references.

L2VPN

This section describes the L2VPN requirements you need to create and configure secure L2VPN connections and identify various supported L2VPN peers.

L2VPN Use Cases

The use cases for Layer-2 VPN are as follows:

- Extend networks (VLANs or VNIs) across sites on the same broadcast domain (with the same subnet)
- Enable hybrid cloud cases
- Enable VM mobility use cases, such as:
 - Migration (including long-distance vSphere vMotion)
 - Disaster recovery without changing the IP address

NSX-T L2VPN

Layer-2 VPN uses GRE over IPsec as a transport method. The edge transport node managed by NSX-T can also be configured as a L2VPN client. The hub-and-spoke topology is supported. Tunnel redundancy is not supported, and L2VPN relies strictly on the edge transport node's high availability. The L2VPN feature is available only in NSX and does not support third-party interoperability.

NSX-T L2VPN Deployment Considerations

When you deploy L2VPN, you must consider that the NSX L2VPN services are supported on Tier-0 and Tier-1 gateways. The segments can be connected to either Tier-1 or Tier-0 gateways to use L2VPN services. Local

egress optimization is supported and pre-shared key (PSK) authentication is the only available method. The peer code is a Base64-encoded string that is available on the L2VPN server. This peer code contains the full L2VPN configuration.

NSX-T L2VPN Hub-and-Spoke

When you configure the Layer-2 VPN using a hub-and-spoke topology, this enables remote sites (spokes) to communicate through a central site (hub). You can configure up to eight L2VPN clients per L2VPN server and the access control is supported through the centralized hub site. The hub acts as a replicator to other client networks. See Figure 6-18.

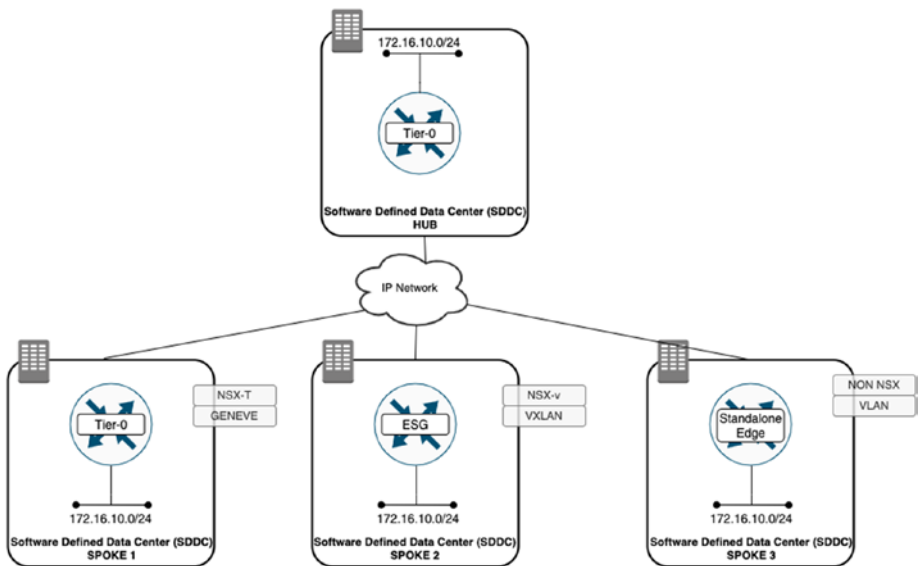


Figure 6-18. Layer-2 VPN in a hub-and-spoke

NSX-T L2VPN Packet Format

The packed format of a Layer-2 tunnel is shown in Figure 6-19.

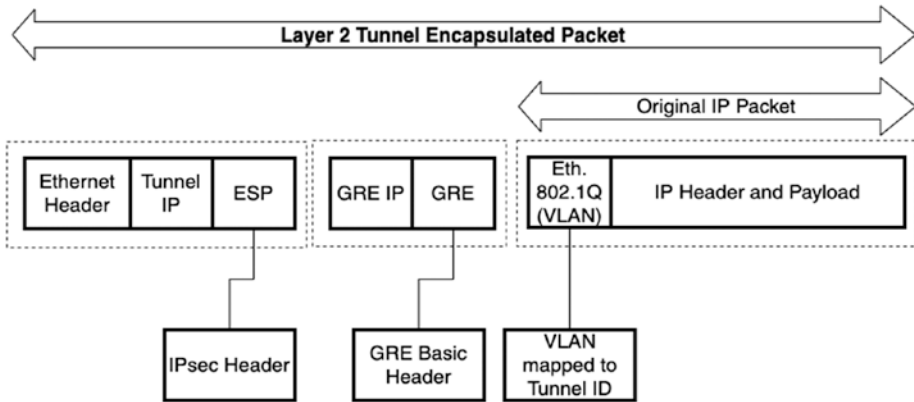


Figure 6-19. Layer-2 tunnel packet format

NSX-T L2VPN Packet Flow

When Layer-2 VPN traffic is sent to a remote destination, the first step is to decapsulate the Geneve frames.

The destination address of the internal frame determines whether traffic goes through the local bridge port toward remote sites or is locally handled.

The other steps involve inserting the ID for the proper VLAN and sending the traffic to the local VTI interface to encapsulate into GRE, which gets protected by IPsec and is forwarded to any given destination.

Then into the inbound direction. When it's receiving L2VPN traffic that is identified by the IPsec engine, the traffic requires IPsec decryption first and GRE decapsulation.

After being sent to the bridge interface, traffic is sent to local networks. The required Geneve encapsulation parameters are based on the actual tunnel IDs for the traffic. See Figure 6-20.

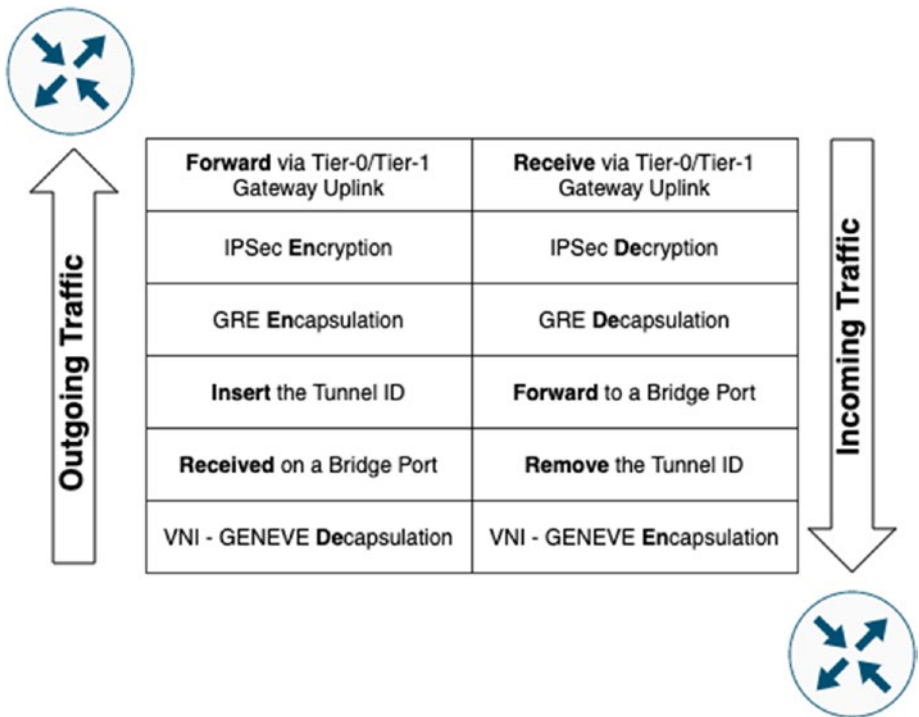


Figure 6-20. Layer-2 VPN traffic packet flow

Layer-2 VPN Scalability

Table 6-13 lists the different sizes of NSX-T Edge transport nodes and the number of VPN sessions that are supported per size.

The numbers are soft limits or tested maximums but may change depending on the NSX-T version. More on the maximums can be found at <https://configmax.vmware.com/>.

Table 6-13. Edge Transport Node L2VPN Session Support Numbers

	Small	Medium	Large	Extra Large	Bare Metal
Number of L2VPN client sessions	1	1	1	1	1
Number of L2VPN server sessions	N/A	128	256	256	256
Segments per session	N/A	512	512	512	512

NSX-T L2VPN Configuration Steps

The L2VPN configuration uses the IPsec features of the gateway. To enable L2VPN, you must first configure an IPsec VPN service and an IPsec local endpoint. Then you can configure the L2VPN details.

When you want to configure a Layer-2 VPN tunnel, you need to perform the configuration steps in [Table 6-14](#).

Table 6-14. L2VPN Configuration Steps

Step	Description
1	Configure the L2VPN server.
2	Configure the L2VPN session.
3	Configure the L2VPN segments.

IPsec Local Endpoint Configuration

You can create a local endpoint by navigating to Networking ► Network Services ► VPN ► Local Endpoints ► Add Local Endpoint. See [Figure 6-21](#).

The screenshot shows the 'ADD LOCAL ENDPOINT' configuration page in the NSX-T VPN interface. At the top, there are navigation tabs: VPN SERVICES, IPSEC SESSIONS, L2 VPN SESSIONS, LOCAL ENDPOINTS (selected), and PROFILES. Below the tabs, there's a 'COLLAPSE ALL' button and a search filter 'Filter by Name, Path and more'. The main form area is titled 'LOCAL ENDPOINT' and contains several sections:

- Name:** A text input field.
- VPN Service:** A dropdown menu labeled 'Select IPsec Serv'. *
- IP Address:** A text input field labeled 'Enter IP Address'.
- Site Certificate:** A dropdown menu labeled 'Select Certificate'.
- Sessions:** A text input field labeled 'Enter Local ID'.
- Status:** A text input field.
- Description:** A text input field.
- Trusted CA Certificates:** A dropdown menu labeled 'Select Certificates'.
- Certificate Revocation List:** A dropdown menu labeled 'Select Certificates'.
- Tags:** A tag selection interface showing 'Tag' and 'Scott' with a plus sign. Below it, it says 'Max 30 allowed. Click (+) to add'.

At the bottom left of the form, there are two buttons: 'SAVE' and 'CANCEL'.

Figure 6-21. IPsec local endpoint (1)

When you configure a new local endpoint, you can specify the parameters outlined in Table 6-15.

Table 6-15. IPsec Local Endpoint Configuration Parameters

Parameter	Description
Name	A name to identify the service.
VPN service	This specifies which IPsec VPN service to use with the endpoint.
IP address	This is for the local IP address.
Site certificate	You use this for certificate-based authentication to specify which certificate to use with this endpoint.
Local ID	This specifies the IPsec ID of the local side. The local ID is usually the same as the local IP address.
Tags	A user-defined tag for references.

NSX-T L2VPN Server Configuration

You can create a Layer-2 VPN server by navigating to Networking ► Network Services ► VPN ► VPN Services ► Add Service: L2VPN Server. See Figures 6-22 and 6-23.




Figure 6-22. Layer-2 VPN server (1)

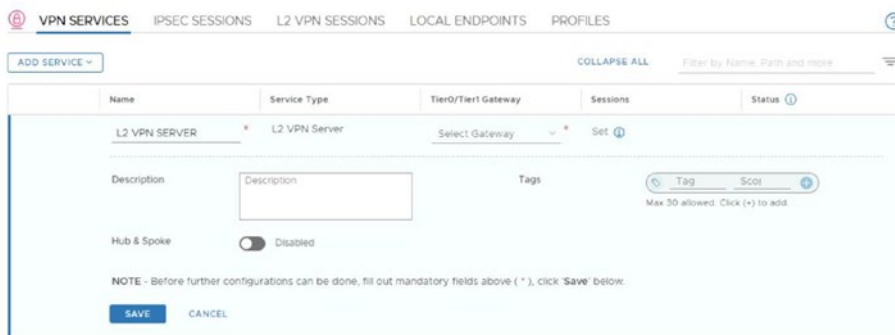


Figure 6-23. Layer-2 VPN server (2)


When you configure a new local endpoint, you can specify the parameters outlined in Table 6-16.

Table 6-16. Local Endpoint Configuration Parameters

Parameter	Description
Name	A name to identify the service.
Tier-0/Tier-1 gateway	This specifies which gateway to use with the endpoint.
Hub & spoke	This enables hub-and-spoke connectivity.
Tag	A user-defined tag for references.

NSX-T L2VPN Session Configuration

You can create a Layer-2 VPN session by navigating to Networking ► Network Services ► VPN ► L2VPN Sessions ► Add L2VPN Session: L2VPN Server. See Figures 6-24 and 6-25.

**Figure 6-24.** Layer-2 VPN session (1)

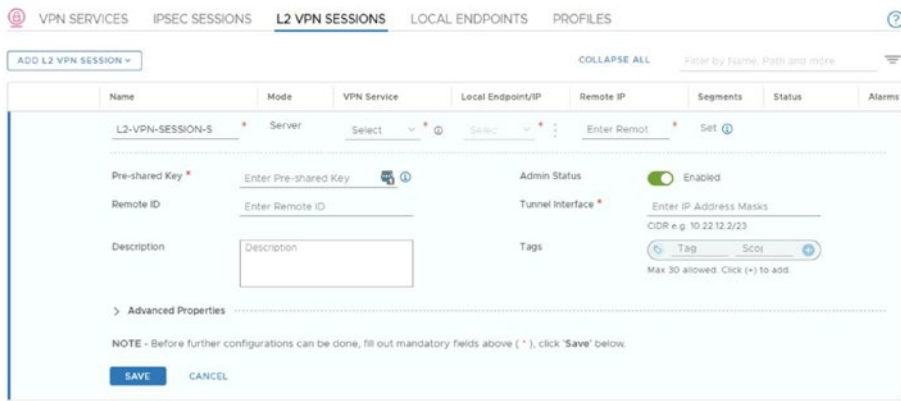


Figure 6-25. Layer-2 VPN session (2)

When you configure a new L2VPN session, you can specify the parameters outlined in Table 6-17.

Table 6-17. Layer-2 VPN Session Configuration Parameters

Parameter	Description
Name	A name to identify the service.
Mode	This is already selected (server).
VPN service	You select the L2VPN service to use.
Local endpoint/IP	The local endpoint configured for the IPsec connection.
Remote IP	The remote IP for the IPsec connection.
Pre-shared key	This is the key or password for this session.
Tunnel interface	You select the IP address of the VTI to use with this session.
Remote ID	This designates the IPsec identifier of the remote IPsec gateway.
Admin status	This is for enabling or disabling the profile.
Tags	A user-defined tag for references.

NSX-T L2VPN Segment Configuration

You identify the segments that should be extended through the L2VPN tunnels. You can edit an existing segment or create a new one.

You can optionally configure a local egress gateway IP so that all VMs on the segment use it as the default gateway.

You can configure a Layer-2 VPN on a segment by navigating to **Networking** ► **Connectivity** ► **Segments** ► **Segments** ► **Add/Select a Segment**. See Figure 6-26.

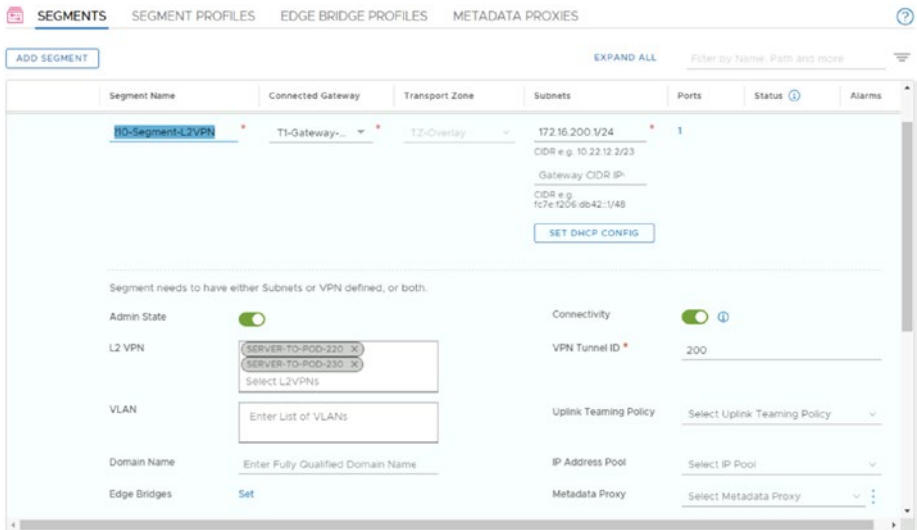


Figure 6-26. Layer-2 VPN on a Segment

When you alter a segment for a L2VPN, you can specify the parameters outlined in Table 6-18.

Table 6-18. Layer-2 VPN Configuration Parameters

Parameter	Description
L2VPN	This defines the previously configured L2VPN session. The segment that is defined is used through that session.
VPN tunnel ID	This number is used to identify the communicating local and remote L2 networks. The same ID on both sides means that they are on the same L2 broadcast domain.
Local egress gateway IP	The IP address of the local gateway that the VMs on the segment use as their default gateway. The same IP address can be configured in the remote site on the extended segment.

You can also configure segments on the L2VPN session page. After selecting the edit mode and clicking Segments, you can add segments and define the tunnel ID for each segment. See Figures 6-27 and 6-28.

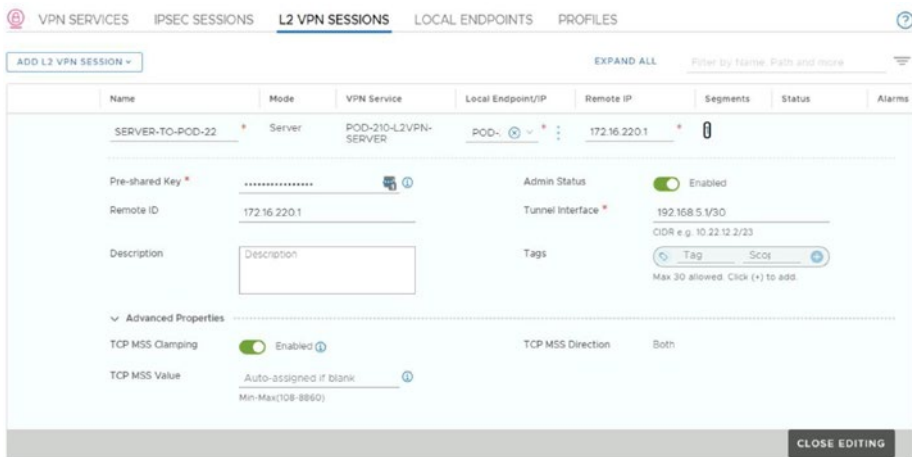


Figure 6-27. L2VPN session (1)



Figure 6-28. L2VPN session (2)

NSX-T L2VPN Supported Clients

You can use different edge options to be configured as L2VPN clients. The L2VPN clients that are supported are listed in Table 6-19.

Table 6-19. L2VPN Client Edge Options

Edge Type	Description
Autonomous edge	The NSX-T autonomous edge instance is deployed by using an OVF file on a host that is not managed by NSX-T.
Standalone edge	The L2VPN client is deployed and configured in an environment that is not managed by NSX. It can be used as an L2VPN client in NSX-v and NSX-T environments. Only VLANs are supported to stretch with this option.
Edge managed by NSX	<p>NSX-v: The edge services gateway in the NSX for vSphere (NSX-v) environment can act as an L2VPN client or server endpoint for an L2VPN.</p> <p>NSX-T: The L2VPN client is deployed and configured in an environment that is managed by NSX-T.</p>

NSX-T L2VPN Peer Compatibility Matrix

Table 6-20 shows the compatibility between the different L2VPN appliances and protocols.

Table 6-20. *L2VPN Appliance Compatibility*

Client/Edge Type	Protocol	NSX-T (3.x)	Network Type
Autonomous edge	L2T	Yes	VLAN
Standalone edge*	SSL	No	VLAN
Standalone edge*	L2T	Yes	VLAN, VXLAN
NSX-managed edge (vSphere)	L2T	Yes	VLAN, VXLAN
NSX-managed edge (NSX-T)	L2T	Yes	VLAN, Geneve
Client	Protocol	NSX-T (3.x)	Network Type

*NSX standalone edge is available for download for both NSX-v and NSX-T.

NSX-T Autonomous Edge

The NSX-T autonomous edge is deployed by using an OVF file on a host that is not managed by NSX-T. The edge is used as an L2VPN client and it acts as an NSX Edge gateway that can be deployed on on-premises data centers and public clouds (for example, Amazon AWS and Microsoft Azure).

The edge runs independently, without being controlled or managed by the management plane/central control plane that is installed in the NSX domain.

The edge can be configured by using REST API and is equipped with the high-performing Data Plane Development Kit (DPDK).

NSX-v Standalone Edge

A standalone edge can be configured as an L2VPN client. A standalone edge must be deployed as a virtual machine (hardware version 11) in a vSphere (6.0 or later) environment.

From the deployment bundle, you select the L2 OVF file (large or extra large) to install an edge VM. During deployment, the L2T section requests the peer code. This code is generated from the server side and should be downloaded/copied and passed into the deployment window.

NSX-v Managed Edge

In NSX data center for vSphere (NSX-v), you can configure the managed edge as an L2VPN client. Starting with NSX-v version 6.4.2, the IPsec-based L2VPN client or server is available on the NSX data center for vSphere edges. The IPsec configuration must be present to make this work; the configuration can only be done through REST API and not from the UI or command line.

NSX-T L2VPN Client Configuration

L2VPN client configuration is similar to the server configuration.

You can create a Layer-2 VPN client by navigating to Networking ► Network Services ► VPN ► VPN Services ► Add L2VPN Session: L2VPN Client. See Figure 6-29.

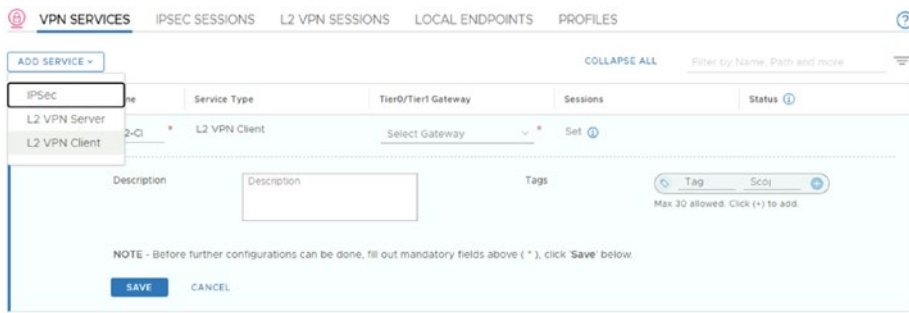


Figure 6-29. L2VPN client

NSX-T L2VPN Session Configuration

Before you create a L2VPN client session, you must get the peer code from the L2VPN server.

The peer code is a Base64-encoded configuration string that is available from the L2VPN server through the DOWNLOAD CONFIG option or through a REST API call.

The REST API call can be found in the NSX-T REST API Reference at <https://code.vmware.com/apis/547/nsx-t>.

You can create a Layer-2 VPN session by navigating to Networking ► Network Services ► VPN ► L2VPN Sessions ► Add L2VPN Session: L2VPN Client. See Figures 6-30 and 6-31.



Figure 6-30. L2VPN session

```

1  {
2  {
3    "transport_tunnel_path":
4    "*/infra/cidr-1a/predefined_id/locale-services/default/ipsec-vpn-services/POD-210-IPSEC/sessions/L2VPN-POD-210-L2VPN-SERVER-5
SESSION",
5    "peer_code":
6    "M0w9ZTQ3ZTAsLHs1c210ZU5hbWU01JNM1ZQT11TRVJWRVtVEStUE9ELTIyMC1eInN0Y1Rho1w1j01MTY5Lj11NC42NC4yI4w1ZHN0VGFwSWA101Ikhjku8tJU
0LjY0LjE1L0JpazVFcHRpb24iOi0pa2V2M1IsImVuY2FwOHJvdG8iOi0ncmVvaXBzIWh1L0JkaEYyb3Vw1j01ZGp4NC1eInN0Y1Z35cHRBbmREeWd1c3Q101Jh2XN
tZ2NtL3N0Y3VyNT11L0Jw2e01017Wkdhcm0iXVZlbnR1eW1wOHVubmVacyI6W3981bG95jYwKzJCi61jE3M14kN14yMjAuMS1eImxvZ2F0Y291c3R5I011eX011uMTY4LjUuM10zNC03N000="
7  }
8  }

```

Figure 6-31. Base64-encoded peer code

NSX-T L2VPN Segment Configuration

You can specify the segments that you want to extend by setting them in the Segments section. See Figure 6-32.

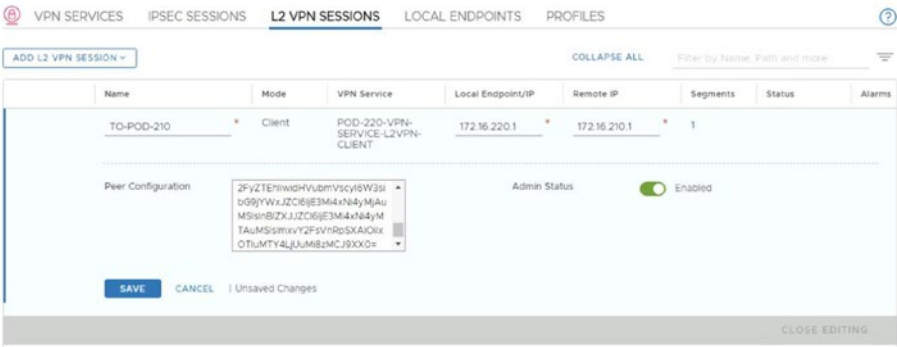


Figure 6-32. L2VPN segment (1)

When you open the Segments area, this is the place where you can add the segment. See Figure 6-33.

Segments - TO-POD-210

L2 VPN Session TO-POD-210 #Segments 1

ADD SEGMENT Search

Segment	VPN Tunnel ID	Local Egress Gateway IP
Select Segment *	Tunnel ID *	Enter IP Address IPv4 Address e.g. 10.10.10.10
Pod-220-Segment-L2VPN	200	

SAVE CANCEL

Figure 6-33. L2VPN segment (2)

I wrote a wiki article on how to configure a Layer-2 hub-and-spoke VPN architecture, which you can find at https://nsx.ninja/index.php/Hub_and_Spoke_Layer_2_VPNs_between_multiple_NSX-T_enabled_sites.

Summary

This chapter explained the differences between an (Layer-3) IPsec VPN and a Layer-2 VPN. It also covered the use cases for each VPN type and explained how they are configured.

The next chapter covers NSX intelligence and discusses other ways to perform operational tasks on NSX-T.

CHAPTER 7

NSX Intelligence, NSX-T Alarms/ Events, and Network Visualizations

This chapter describes NSX Intelligence and its use cases. It explains the NSX Intelligence system requirements and discusses what you need to deploy the NSX Intelligence appliance. It also shows you how NSX Intelligence performs visualization and recommendation capabilities and describes NSX Intelligence alarms and events so you can identify and prevent failures of your NSX-T environment. The chapter also describes the different use cases for the network topology feature.

NSX Intelligence

NSX Intelligence is a distributed “big-data” analytics tool that provides visibility and dynamic security policy enforcement for NSX-T environments.

It allows you to visualize data center objects and traffic flows and provides you with recommendations for micro-segmentation firewall policies, including the distributed firewall rule recommendation with the continuous monitoring of changes in the environment. It also collects alarms and events so you can identify and prevent failures of your NSX-T environment.

The NSX Intelligence appliance is a separate appliance that needs to be deployed from the NSX-T GUI and will integrate during the deployment with your NSX environment.

NSX Intelligence Use Cases

NSX Intelligence has three clearly defined use cases:

- Provides visibility in an NSX domain
- Provides (micro-segmentation) security policy recommendations of network traffic that flows between the virtual machines that are “networked” with NSX-T.
- Provides constant traffic evaluation and analysis

NSX Intelligence vs vRealize Network Insight

NSX Intelligence and vRealize Network Insight are two products that may be presumed to have similar functionalities. However, there are some differences. They complement each other to deliver better security and network operations.

NSX Intelligence is typically used for micro-segmentation planning and vRealize Network Insight is used for virtual and physical network monitoring. Table 7-1 explains these differences in more detail.

Table 7-1. *NSX Intelligence vs vRealize Network Insight*

NSX Intelligence for Security Operations	vRealize Network Insight for Network Operations
Gains insight about the east-west traffic flows.	Provides end-to-end network visibility for physical, virtual, VMware SD-WAN, cloud, and third-party environments.
Simplifies the (DFW) rule recommendation and deployment on NSX-T.	Provides day-2 network operations and troubleshooting

NSX Intelligence Requirements

You must deploy NSX Intelligence on an ESXi host that is managed by a vCenter Server that's registered as a compute manager in NSX Manager. You need to allocate an IP address for NSX Intelligence, as the appliance must have a static IP address. You cannot change the IP address after installation.

NSX Intelligence can only be deployed from an NSX Manager (standalone or cluster). If you're using NSX Federation, you will not find this option in the Global Manager UI. To use NSX Intelligence, you need to have an NSX-T Data Center Enterprise Plus license. You must also have an Enterprise Administrator role to install, configure, and use NSX Intelligence.

NSX Intelligence Footprint

NSX Intelligence is pretty “resource hungry” because it performs constant analytics on your NSX-T environment. There are two deployment sizes you can pick from, as outlined in Table 7-2.

Table 7-2. *NSX Intelligence Deployment Sizes*

	Small Appliance	Large Appliance
Environment	Can be used in a lab or proof-of-concept environment or in a small-scale production environment.	Can be used in a large-scale production environment.
vCPU	16	32
RAM	64GB	28GB
Disk	2TB	2TB

NSX Intelligence Deployment

In the NSX-T Manager, navigate to Plan & Troubleshoot ► Discover & Take Action.

You will see the screen in Figure 7-1. Click Go To System to start the deployment of the NSX Intelligence appliance.

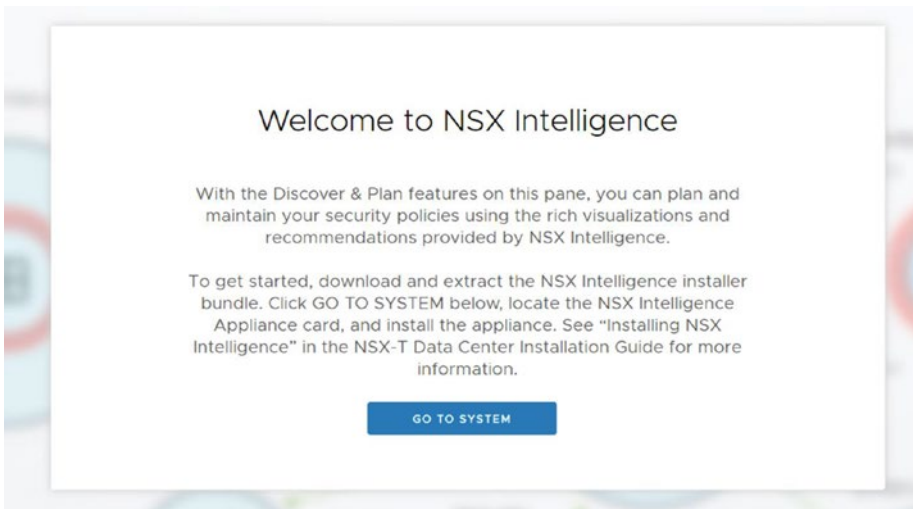


Figure 7-1. *NSX Intelligence welcome screen (before deployment)*

You can deploy the NSX Intelligence appliance from the NSX UI by manually navigating to System ► Configuration ► Appliances ► NSX Intelligence ► Add NSX Intelligence Appliance. See Figure 7-2.

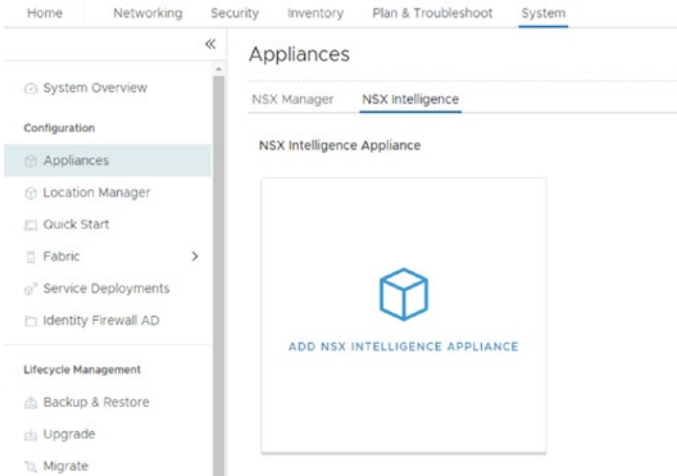


Figure 7-2. Deploy the NSX Intelligence appliance (1)

First you need to upload the .ova of the NSX Intelligence appliance to the NSX-T Manager node. This is currently done with the option to select the .ova file from the deployment screen. See Figure 7-3.

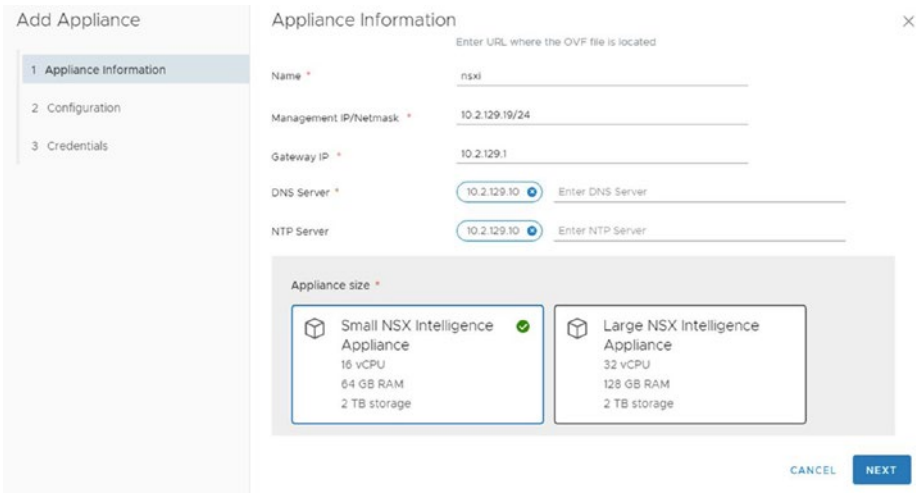


Figure 7-3. Deploy the NSX Intelligence appliance (2)

In the next step, you configure the vSphere details for the virtual appliance. See Figure 7-4.

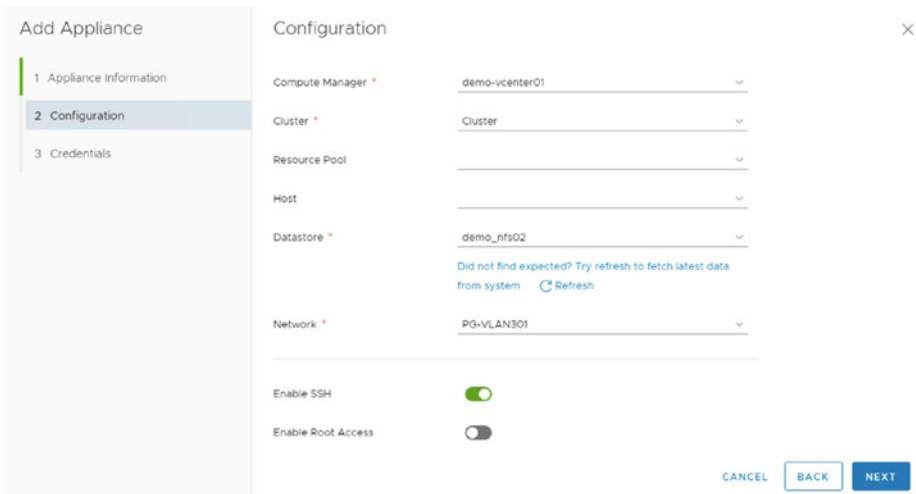


Figure 7-4. Deploy the NSX Intelligence appliance (3)

You then configure the appliance credentials during the third and final step. See Figure 7-5.

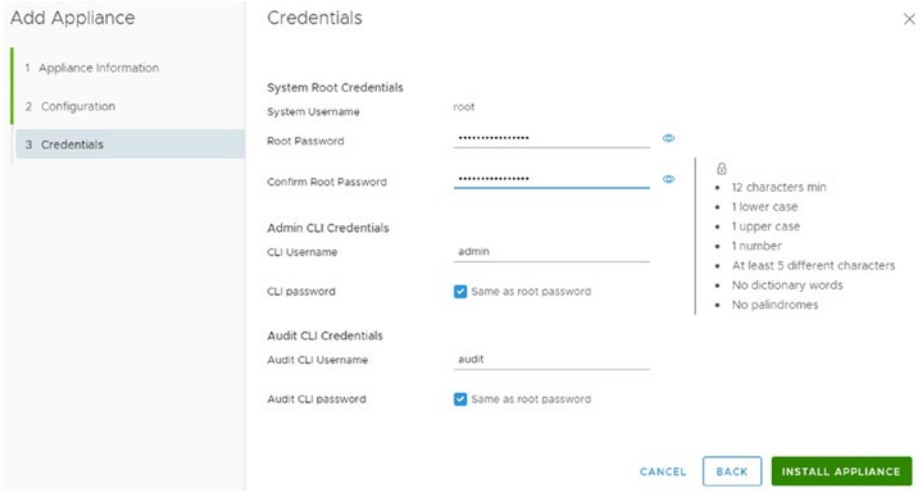


Figure 7-5. Deploy the NSX Intelligence appliance (4)

Click Install Appliance to start the deployment. See Figure 7-6.

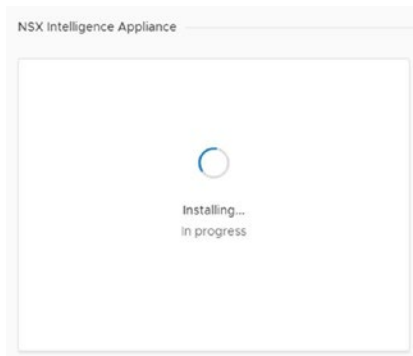


Figure 7-6. Deploy the NSX Intelligence appliance (5)

The deployment will take about five minutes. You'll see the screen in Figure 7-7 when it's been deployed successfully.

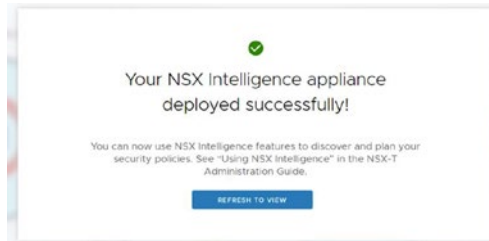


Figure 7-7. Deploy the NSX Intelligence appliance (6)

NSX Intelligence Verification

When you successfully finish the deployment, you can validate the status of the NSX Intelligence appliance by navigating to System ► Configuration ► Appliances ► NSX Intelligence.

Data collection starts automatically after deploying the appliance. You can stop or start data collection from the Actions menu and delete the appliance from there as well.

Flow Visualization with NSX Intelligence

You can visualize flows for virtual machine (objects) and security groups by navigating to Plan & Troubleshoot ► Discover & Plan ► Discover & Take Action.

Although it's a separate virtual appliance, the NSX Intelligence UI seamlessly integrates with the NSX Manager UI. It can be found by choosing Plan & Troubleshoot ► Discover & Take Action. See Figure 7-8.

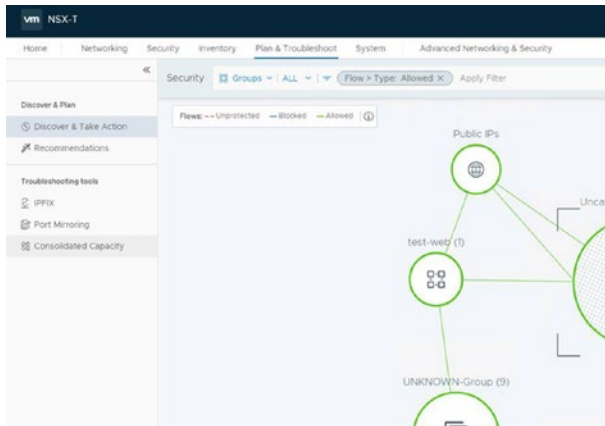


Figure 7-8. Flow visualization

You need to make sure you have created Groups by going to Inventory See Figure 7-2. Groups. NSX Intelligence will use these groups as a starting point to create visualizations.

The two objects you can work with here are virtual machines and groups, as shown in Figure 7-9.

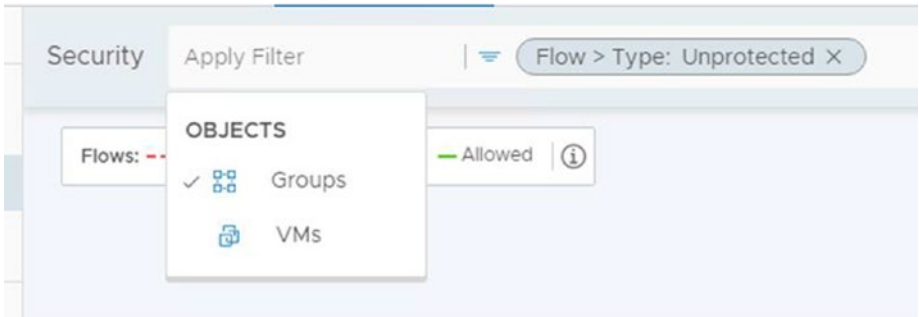


Figure 7-9. Flow visualization based on groups

You can choose to display only certain VMs/groups or all of them, as shown in Figure 7-10.

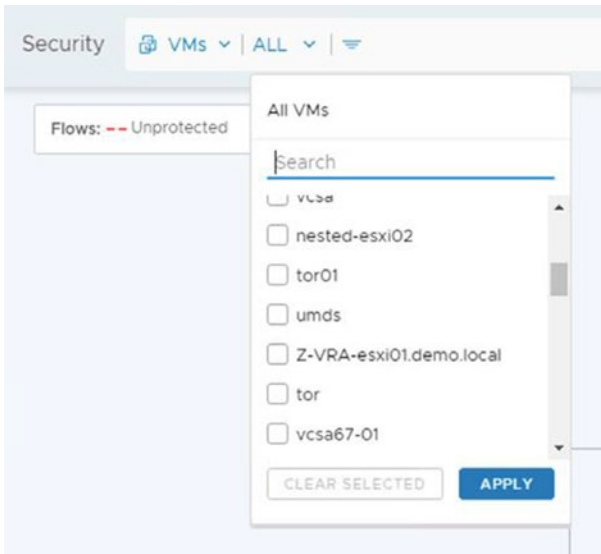


Figure 7-10. Flow visualization based on virtual machines

You can also apply a filter based on tags, flows, and rules. See Figure 7-11.

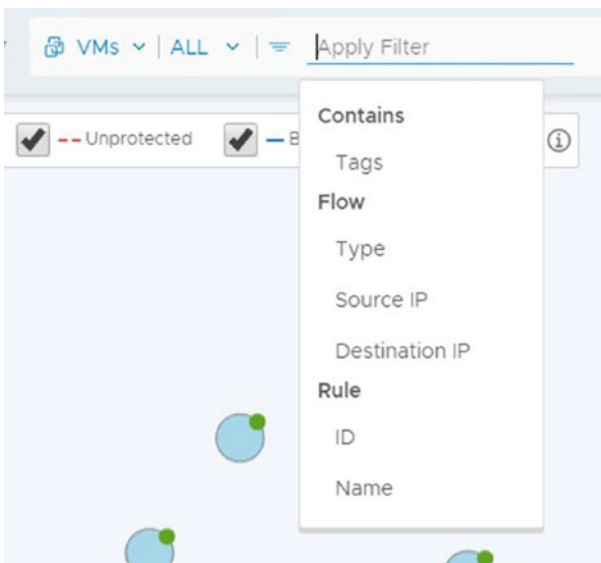


Figure 7-11. Flow visualization based on tags, flows, and rules

When you power on two Windows virtual machines, it will take about 20 seconds before the NSX Intelligence engine starts to draw the communication paths of these virtual machines. See Figure 7-12.

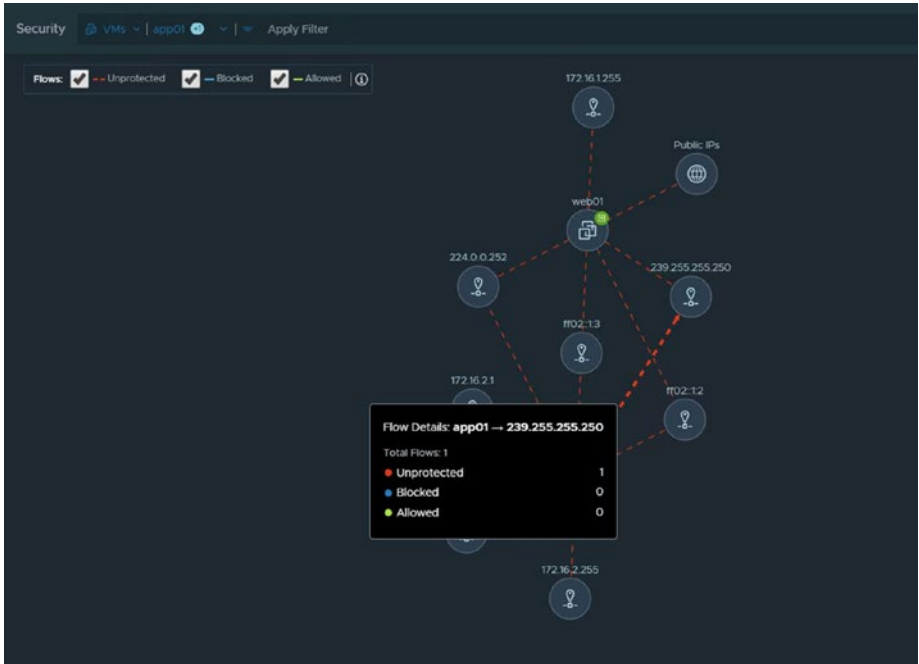


Figure 7-12. Communication path visualization between virtual machines

Recommendations with NSX Intelligence

You can start a new security recommendation by navigating to Plan & Troubleshoot > Discover & Plan > Recommendations. It is possible to run multiple concurrent recommendations at the same time.

To get actual firewall rule recommendations, you need to start a new recommendation process. See Figure 7-13.

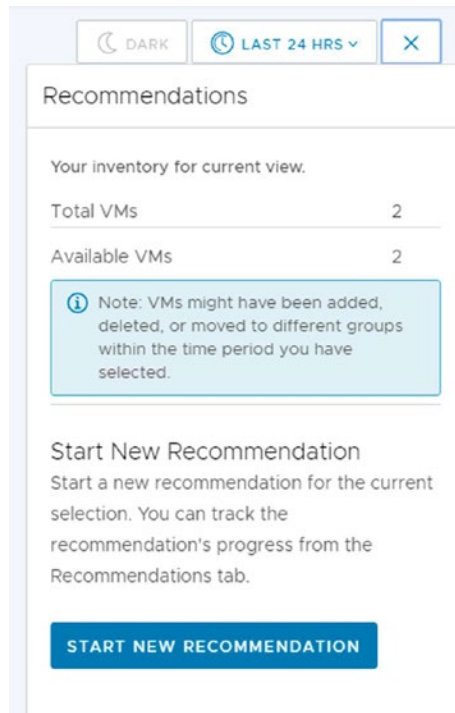


Figure 7-13. Firewall rule recommendations (1)

After clicking the Start Recommendation button, you can configure some parameters, Time Range being the most important. See Figure 7-14.

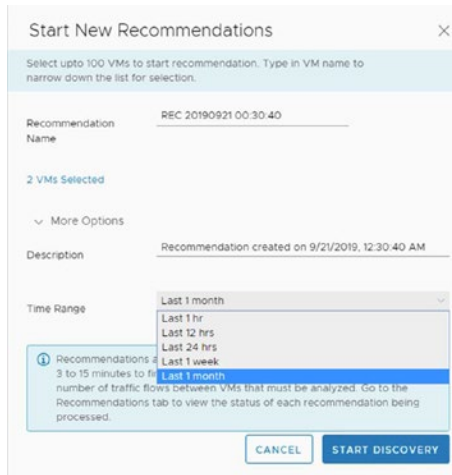


Figure 7-14. Firewall rule recommendations (2)

Click Start Discovery to kick off the recommendation process. This process can be monitored under Recommendations, as shown in Figure 7-15.




Figure 7-15. Firewall rule recommendations (3)

Once the analysis of the recommended rules is complete, then groups and services can be reviewed and modified. See Figure 7-16.

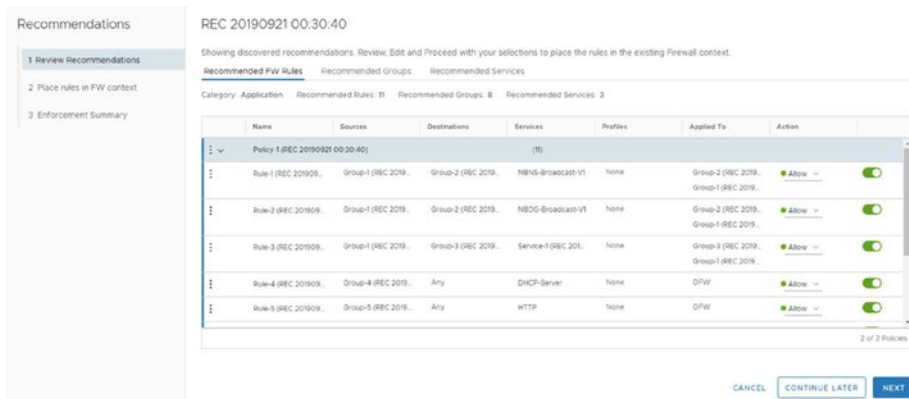


Figure 7-16. Firewall rule recommendations (4)

In Step 2, you can choose placement for the new recommendations based the security policies. See Figure 7-17.

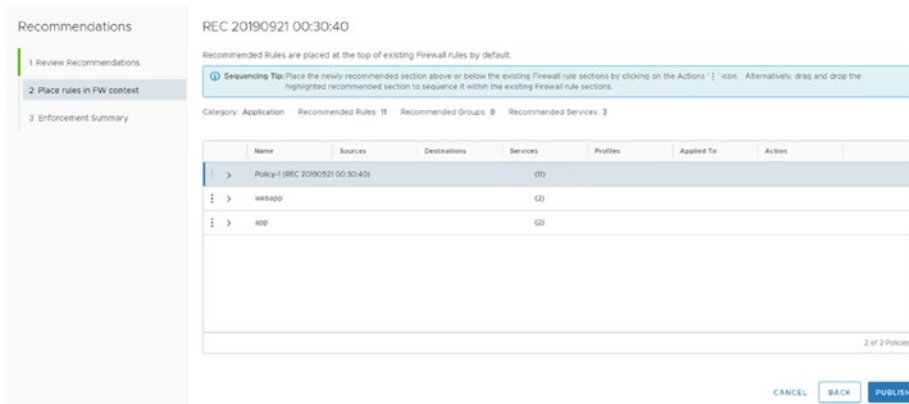


Figure 7-17. Firewall rule recommendations (5)

When you click Publish, NSX-T will create the objects and enforce the security policy in the distributed firewall. See Figure 7-18.



Recommended security policies have been successfully published.

To view and manage published recommended rules, go to **Security -> Distributed Firewall**

Figure 7-18. Firewall rule recommendations (6)

You can verify that the recommended rules are in place when you browse to the Distributed Firewall configuration page. See Figure 7-19.

<input type="checkbox"/>	Name	Sources	Destinations	Services	Profiles	Applied To	Action	
▼	Policy-1 (REC 20190... (7))	Applied To: DFW						● Up ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-1 (REC 201909...)	Group-1 (REC 201...)	Group-2 (REC 201...)	NBNS-Broadcast...	None	Group-2 (REC 201...) Group-1 (REC 201...)	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-2 (REC 20190...)	Group-1 (REC 201...)	Group-2 (REC 201...)	NBDG-Broadcast...	None	Group-2 (REC 201...) Group-1 (REC 201...)	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-3 (REC 20190...)	Group-1 (REC 201...)	Group-3 (REC 201...)	Service-1 (REC 20...	None	Group-3 (REC 201...) Group-1 (REC 201...)	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-4 (REC 20190...)	Group-4 (REC 201...)	Any	DHCP-Server	None	DFW	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-5 (REC 20190...)	Group-5 (REC 201...)	Any	HTTP	None	DFW	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-6 (REC 20190...)	Group-5 (REC 201...)	Any	Service-2 (REC 2...	None	DFW	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-7 (REC 20190...)	Group-5 (REC 201...)	Any	HTTPS	None	DFW	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-8 (REC 20190...)	Group-5 (REC 201...)	Any	NTP	None	DFW	● Allow ▾	● ⌵ ⌶ ⌷
⋮	<input type="checkbox"/> Rule-9 (REC 20190...)	Group-1 (REC 201...)	Group-5 (REC 201...)	Service-3 (REC 2...	None	Group-6 (REC 201...) Group-1 (REC 201...)	● Allow ▾	● ⌵ ⌶ ⌷

Figure 7-19. Firewall rule recommendations (7)

Alarms with NSX Intelligence

NSX Intelligence triggers alarms in the NSX GUI based on several events:

- CPU usage
- Memory usage
- Disk and data disk partition usage
- Node status

NSX-T Network Topology

Using the NSX-T Network Topology feature, NSX administrators can view the configured network. You can view all Tier-0 and Tier-1 gateways with the attached segments and workloads.

NSX-T Network Topology Use Cases

By using the network topology, you can visualize how the different components are interconnected and gain a better understanding of your network topology. You can view the VMs connected to a segment and view the details of the virtual machines, segments, and Tier-1 and Tier-0 gateways with the IP addresses that are configured in the network.

NSX-T Network Topology GUI

Figures 7-20 through 7-22 display the options that are available to control the network topology view. You can click or point to the virtual machines, segments, and Tier-1 and Tier-0 gateways to display more information.

When you zoom out, you can see the topology from a very high level, without any detailed information. See Figure 7-20.




Figure 7-20. Network topology view (1)

When you zoom in, you can see more details like the names and IP addresses. See Figure 7-21.

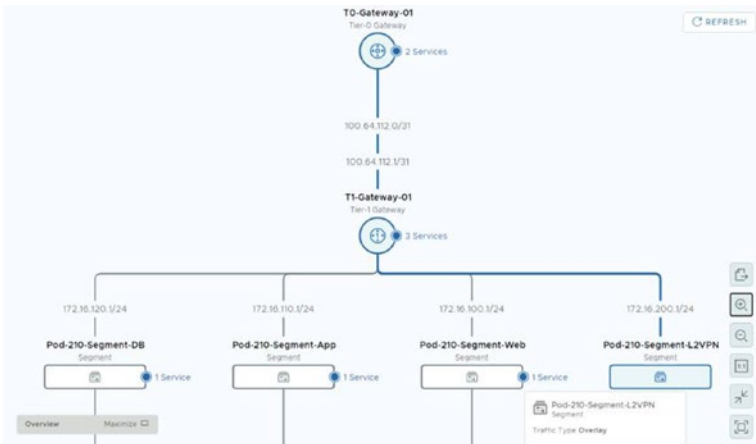


Figure 7-21. Network topology view (2)

You can also click certain objects to see more object-related information, such as the Tier-0 gateway. See Figure 7-22.

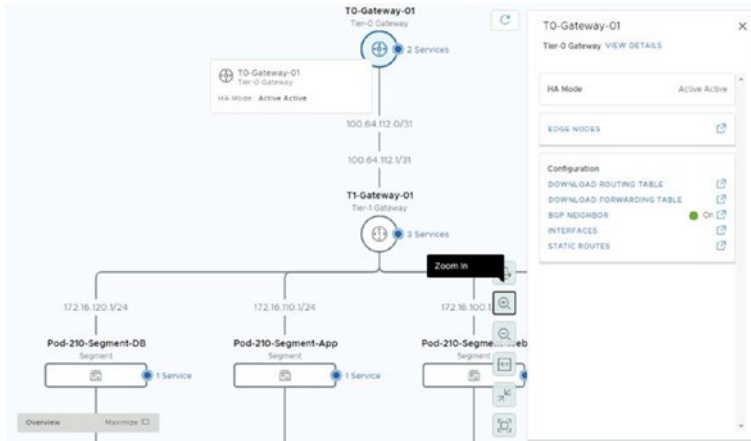


Figure 7-22. Network topology view (3)

NSX-T Alarms and Events

This section explains how to define alarms and events. It describes the alarms and events terminology and describes the architecture of alarms.

NSX-T Native Monitoring Tools

NSX-T offers a variety of tools to monitor different things. Table 7-3 shows these different monitoring tools.

Table 7-3. *NSX-T Monitoring Tools*

NSX-T Tool	Purpose	Autonotification
Dashboard	Monitoring	N/A
Logs	Troubleshooting	N/A
CLI	Troubleshooting	N/A
API	Configuration troubleshooting	N/A
SNMP	Monitoring notification	Yes (traps and polling)
Alarms and Events	Monitoring troubleshooting	Yes

NSX-T Events

NSX-T events are records of user-generated or system-generated actions that occur in NSX-T Manager or on a transport node. The events data includes details about the event, such as who generated the event, the type of the event, and when and where the event occurred.

NSX-T Alarms

NSX-T alarms are notifications that are activated in response to an event, a set of conditions, or the state of an object in the NSX-T system. An alarm is uniquely identified by its UUID and the alarm severity levels are Critical, High, Medium, or Low.

NSX-T Alarm Use Cases

Alarms can be used to monitor and troubleshoot the NSX-T Management, control and data plane.

You can also monitor and troubleshoot NSX-T services, such as segments, Tier-1 and Tier-0 gateways, and load balancers.

NSX-T Alarm and Events Architecture

Figure 7-23 shows how the NSX-T Manager collects the alarms and events from the transport nodes:

1. The client library tracks events on the transport nodes and passes them to the nsx-opsAgent.
2. The nsx-opsAgent process forwards the events to the nsx-proxy.
3. The nsx-proxy sends the event to the appliance proxy hub.
4. The manager service receives events from registered nodes.
5. The manager writes the alarms in the Corfu database.

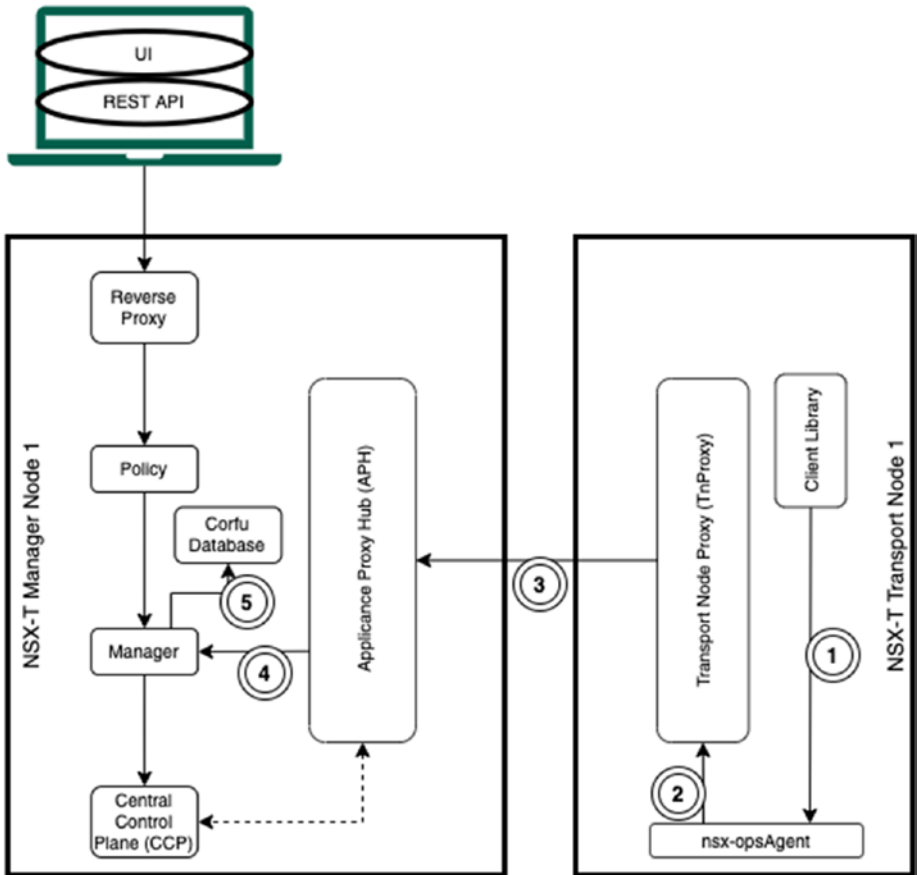


Figure 7-23. NSX-T alarm and events architecture

NSX-T Alarm Definitions

You can see the different alarm definitions in the NSX-T GUI by navigating to Home ► Alarms ► Alarm Definitions. See Figure 7-24.

Feature	Event Type	Severity	Enabled	Create Alarms	Create SNMP Traps
Alarm Management	Alarm Service Overloaded	Critical	Yes	Yes	Yes
Alarm Management	Heavy Volume Of Alarms	Critical	Yes	Yes	Yes
Audit Log Health	Audit Log File Update Error	Critical	Yes	Yes	Yes
Audit Log Health	Remote Logging Server Error	Critical	Yes	Yes	Yes
Capacity	Maximum Capacity	Critical	Yes	Yes	Yes
Capacity	Maximum Capacity Threshold	High	Yes	Yes	Yes
Capacity	Minimum Capacity Threshold	Medium	Yes	Yes	Yes
Certificates	Certificate Expiration Approaching	Medium	Yes	Yes	Yes
Certificates	Certificate Expired	Critical	Yes	Yes	Yes
Certificates	Certificate is About To Expire	High	Yes	Yes	Yes
Cx Health	Hyperbus Manager Connection Down	Medium	Yes	Yes	Yes
Communication	Control Channel To Manager Node Down	Medium	Yes	Yes	Yes
Communication	Control Channel To Manager Node Down Too Long	Critical	Yes	Yes	Yes

Figure 7-24. NSX-T Alarm Definitions

NSX-T Alarm Definition Configuration

You can customize specific alarm definitions by changing their threshold values, sensitivity, and much more by editing the specific alarm definition. See Figure 7-25.

Feature	Event Type	Severity	Enabled	Create Alarms	Create SNMP Traps
<input type="checkbox"/> Edit	Management	Alarm Service Overloaded	Critical	Yes	Yes

Figure 7-25. NSX-T alarm definition configuration (1)

Different alarm definitions will provide you with different configurable items. See Figure 7-26.




Figure 7-26. NSX-T alarm definition configuration (2)

NSX-T Alarm Verification

You can verify existing alarms from the NSX GUI by navigating to Home ► Alarms ► Alarms. See Figure 7-27.

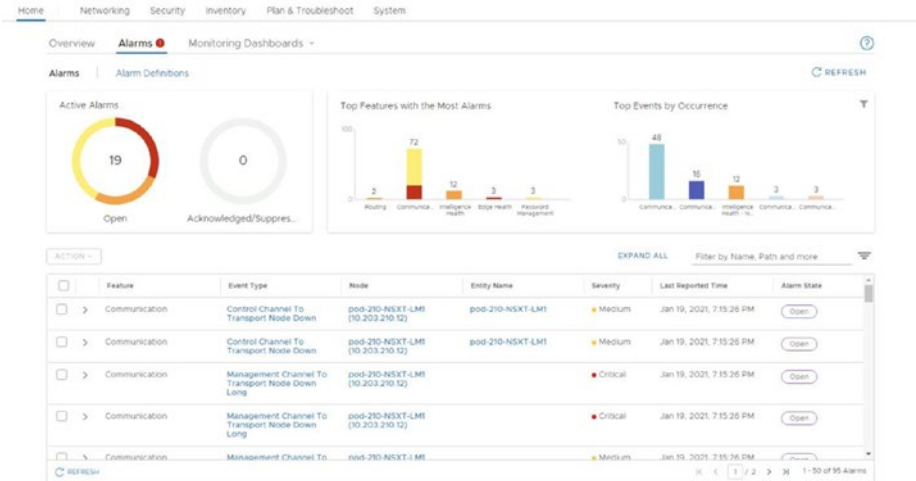


Figure 7-27. NSX-T alarm verification

NSX-T Alarm Actions

When you select a specific alarm, you can acknowledge, resolve, suppress, or open the alarm message. See Figure 7-28.

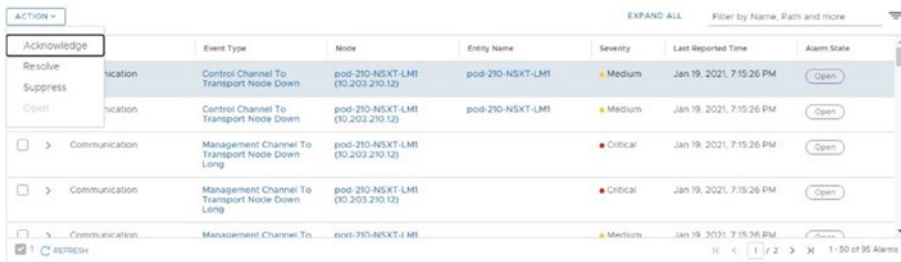


Figure 7-28. The NSX-T Acknowledge alarm action

NSX-T Alarm Suppression

If you select suppression, you can suppress the alarm for a specific duration, in hours. See Figures 7-29 through 7-31.



Figure 7-29. NSX-T Suppression alarm action (1)




Figure 7-30. NSX-T Suppression alarm action (2)




Figure 7-31. NSX-T Suppression alarm action (3)

View NSX-T Open Alarms

For some specific objects or NSX-T components, there will be a separate column available that will display their alarms.

The Alarms column displays the number of alarms generated on Tier-0 and Tier-1 gateways, segments, and load balancers.

Figure 7-32 shows this column for the edge transport nodes.




Figure 7-32. NSX-T open alarms

Summary

This chapter covered what NSX Intelligence is and how it's used. It compared NSX Intelligence to vRealize Network Insight. It also discussed what information you can retrieve from the network topology feature and how alarms and events can be used to keep your NSX-T environment healthy.

The next chapter covers authentication and authorization in combination with NSX-T.

CHAPTER 8

Authentication and Authorization

This chapter describes the purpose of VMware Identity Manager (also known as VMware Workspace ONE Access) and identifies the benefits of integrating NSX-T VMware Identity Manager (vIDM). You will learn how to configure the integration between NSX-T and vIDM and how to verify it.

The chapter also describes the LDAP authentication architecture and identifies the benefits of integrating NSX-T with LDAP. It explains how to identify the different types of users and authentication policies available in NSX-T and how to recognize permissions and roles that NSX-T has to offer.

VMware Identity Manager (vIDM)

The VMware Identity Manager (vIDM) is an identity as a service (IDaaS) solution that provides the following services for software as a service (SaaS)—web, cloud, and native mobile applications like application provisioning, conditional access control, and single sign-on (SSO) services.

A variety of VMware products can use vIDM as an enterprise SSO solution; it's based on the OAuth 2.0 authorization framework.

vIDM with NSX-T Integration

You have the option to integrate NSX-T with vIDM and configure Role Based Access Control (RBAC) for users managed by vIDM.

Integration provides a variety of benefits related to user authentication, as listed in Table 8-1.

Table 8-1. *vIDM Integration with NSX-T Benefits*

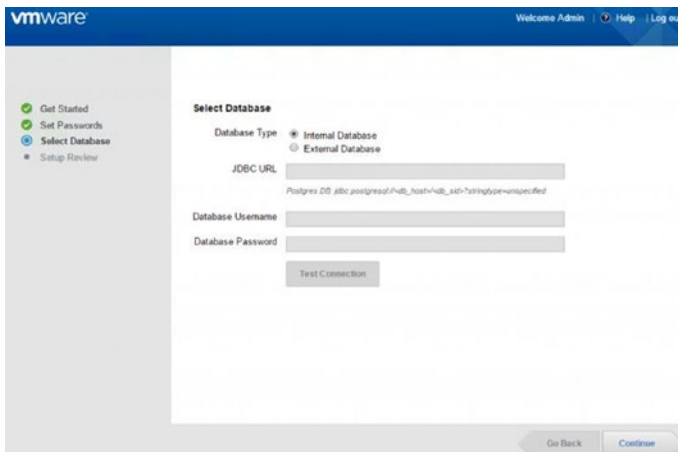
Authentication,	• RADIUS
Authorization, and	• Smart cards and common access cards
Accounting (AAA) systems supported	• RSA SecureID
	• LDAP and OpenLDAP based on Active Directory (AD)
Enterprise SSO	• Common authentication platform across multiple VMware solutions
	• Seamless SSO experience

VIDM with NSX-T integration Prerequisites

Before integrating VMware Identity Manager with NSX-T, you need to complete the steps outlined in Table 8-2. See Figure 8-1.

Table 8-2. *vIDM Integration Prerequisites*

Step	Description
1	Deploy the VMware Identity Manager appliance from an OVF template.
2	<p>Synchronize the VMware Identity Manager virtual machine time with the ESXi host where it is running.</p> <ul style="list-style-type: none"> • Right-click the VM and select Edit Settings ► VM Options. • Scroll down to the Time section and select the Synchronize Guest Time with Host check box. <p>Another option is to use an external NTP server for the time synchronization.</p>
3	<p>After deploying the VMware Identity Manager appliance, use the Setup wizard available at <a href="https://<VMware_Identity_Manager_FQDN>">https://<VMware_Identity_Manager_FQDN> to set passwords for the admin, root, and remote SSH user and to select a database.</p> <p>You can use the external Microsoft SQL or Oracle database or the internal PostgreSQL database.</p>

**Figure 8-1.** *Database selection when deploying vIDM*

VIDM with NSX-T Integration

After initially setting up vIDM, you need to log into the vIDM administration console. In the console, you can configure identity sources, authentication methods, and access policies.

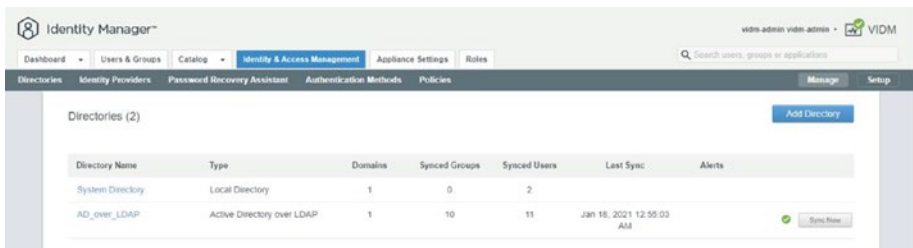
You can access the VMware Identity Manager Administration console via the following URL `https://<VMware_Identity_Manager_FQDN>:443/SAAS/admin`.

The following identity sources are supported in the VMware Identity Manager.

To configure authentication methods, select Identity & Access Management ► Authentication Methods.

Administrators can configure a single authentication method or set up chained, two-factor authentication. To define access policies, select Identity & Access Management ► Policies.

Administrators can configure rules that specify the network ranges and types of devices that users can use to sign in. See Figure 8-2.



Directory Name	Type	Domains	Synced Groups	Synced Users	Last Sync	Alerts
System Directory	Local Directory	1	0	2		
AD_over_LDAP	Active Directory over LDAP	1	10	11	Jan 18, 2021 12:55:03 AM	

Figure 8-2. vIDM authentication methods

When NSX-T and vIDM are both deployed and configured, you can integrate these components with each other by following the steps in Table 8-3.

Table 8-3. *vIDM with NSX-T Integration Steps (vIDM Side)*


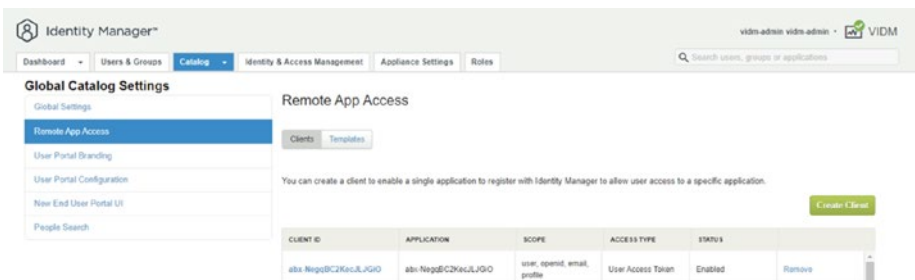
Step	Description
1	Create an OAuth client for NSX in VMware Identity Manager.
2	Obtain the SHA-256 certificate thumbprint (from the NSX-T Manager) that you need to configure on the VMware Identity Manager appliance.
3	Configure VMware Identity Manager details in NSX.

Creating an OAuth Client

vIDM uses the OAuth 2.0 authorization framework to enable third-party applications, such as NSX and its users, to access specific data and services. vIDM protects the third party's account credentials.

Before you can integrate vIDM with an NSX-T data center, you must register NSX-T as a trusted OAuth client in vIDM.

Use the vIDM administration console and click the Catalog tab. Select Settings ► Remote App Access. On the Clients page, click Create Client. See Figures 8-3 and 8-4.

**Figure 8-3.** *Create OAuth client (1)***Figure 8-4.** *Create OAuth client (2)*

When configuring NSX-T data center details, you select Service Client Token from the Access Type drop-down menu. This selection indicates that the application, NSX-T Data Center in this example, accesses the APIs itself. The application does not access the APIs for a particular user.

You must specify a client ID to uniquely identify NSX. You need this value to enable the VMware Identity Manager integration.

You must also click Generate Shared Secret. You need this value to enable the VMware Identity Manager integration. Leave the default settings for all other options.

On the Create Client page, you can optionally set the token time-to-live values by specifying the access, refresh, and idle timers (Figures 8-5 and 8-6).

Create Client

Access Type* Service Client Token

Client ID* pod-100-nsxt-lm1
Characters allowed are: alphanumeric (A-Z, a-z, 0-9) period (.), underscore (_), and hyphen (-) and at sign (@).

Scope* admin

Advanced ▼

Shared Secret amDE8WWUyZIBFRzmSEGRpwVai8FaMBKh2DY91rsIjgDc
[Generate Shared Secret](#)

Issue Refresh Token Refresh Token

Token Type Bearer

Access Token Time-To-Live (TTL) 3 hours

Refresh Token Time-To-Live (TTL) 90 days (1 day is 24 hours)

Idle Token Time-To-Live (TTL) 4 days (1 day is 24 hours)

* Required

Cancel Add

Figure 8-5. Create OAuth client (3)

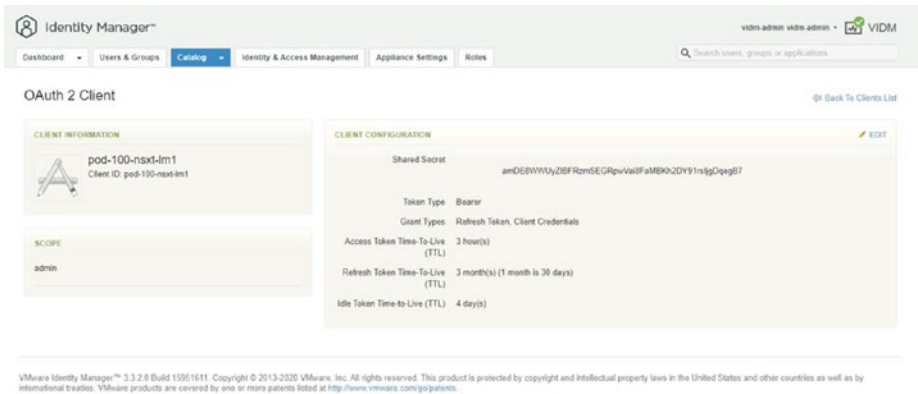


Figure 8-6. Create OAuth client (4)

SHA-256 Certificate Thumbprint

You need the SHA-256 certificate thumbprint value when you enable the vIDM integration.

Before you configure the integration between NSX-T vIDM, you must obtain the SHA-256 certificate thumbprint of the vIDM appliance. Use the steps in Table 8-4 to do so.

Table 8-4. Obtain the SHA-256 Certificate Thumbprint of the vIDM Appliance

Step	Description
1	Use SSH to log into the VMware Identity Manager appliance.
2	Run the <code>sudo -s</code> command to gain root access.
3	Navigate to the VMware Identity Manager configuration directory using the following command: <code>cd /usr/local/horizon/conf</code>
4	Retrieve the SHA-256 certificate thumbprint of vIDM using the following command: <code>openssl x509 -in <vidm_fqdn>.pem -noout -sha256 -fingerprint</code>

For example:

```
vidm:~ # sudo -s
vidm:~ # cd /usr/local/horizon/conf
vidm:/usr/local/horizon/conf # openssl x509 -in vidm.sddc.lab_
cert.pem -noout -sha256 -fingerprint
SHA256 Fingerprint=2A:7E:AC:1C:58:4E:D2:65:39:B0:3E:46:B8:F7:30
:EE:4F:E7:C5:26:B0:08:6D:6C:B1:B8:E8:FE:44:BB:00:9F
vidm:/usr/local/horizon/conf #
```

VIDM with NSX-T Integration Configuration

You can now enable vIDM integration from the NSX GUI (Figure 8-7). Navigate to System ► Settings ► Users and Roles. Then follow the steps outlined in Table 8-5.

Table 8-5. *vIDM with NSX-T Integration Steps (NSX-T Side)*

Steps	Description
1	Click the VMware Identity Manager tab.
2	Click Edit.
3	Provide the configuration information and click Save.

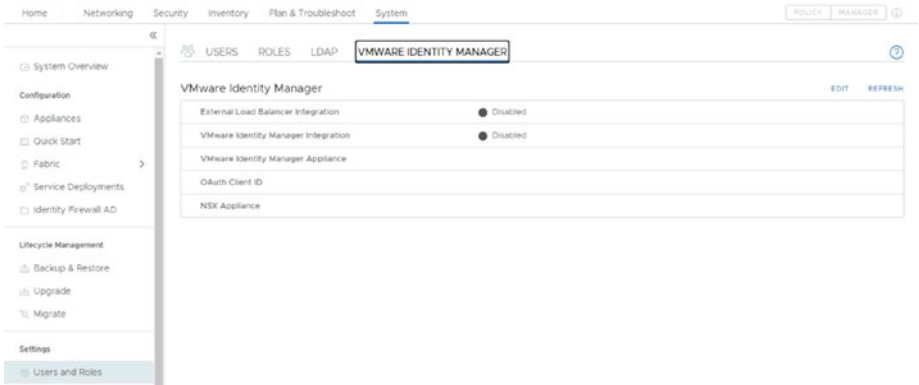


Figure 8-7. vIDM with NSX-T integration (NSX-T side) (1)

In the OAuth Client ID and OAuth Client Secret text boxes, enter the client ID and shared secret that you generated when you created the OAuth client for NSX-T in vIDM.

In the SSL Thumbprint text box, enter the SHA-256 certificate thumbprint value that you generated from the vIDM appliance command line (provided in the previous chapter).

The value entered in the NSX Appliance text box must be used to access the NSX Manager after integration. If you enter the fully qualified domain name (FQDN) of the NSX Manager and try to access the appliance through its IP address, the authentication will fail.

If a virtual IP (VIP) is set up in the NSX Management cluster, you cannot use the external load-balancer integration even if you enable it. You can either have VIP or the external load balancer while configuring VMware Identity Manager, but not both. Disable VIP if you want to use an external load balancer. See Figures 8-8 and 8-9.

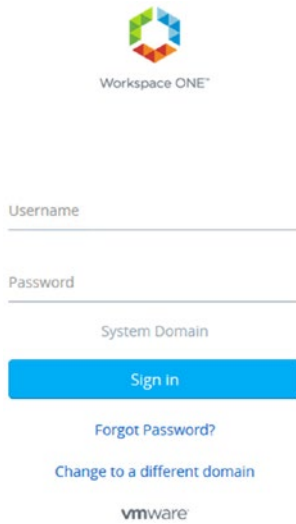


Figure 8-8. vIDM with NSX-T integration (NSX-T side) (2)

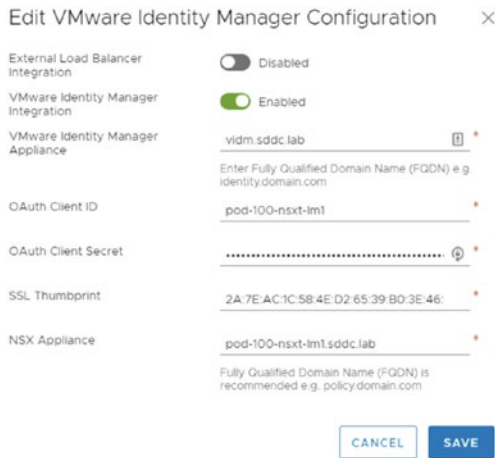


Figure 8-9. vIDM with NSX-T integration (NSX-T side) (3)

vIDM with NSX-T Integration Verification

When you navigate to System ► Settings ► Users and Roles ► VMware Identity Manager, you can validate the vIDM integration.

When the integration is successful, the vIDM connection appears, as shown in Figure 8-10.

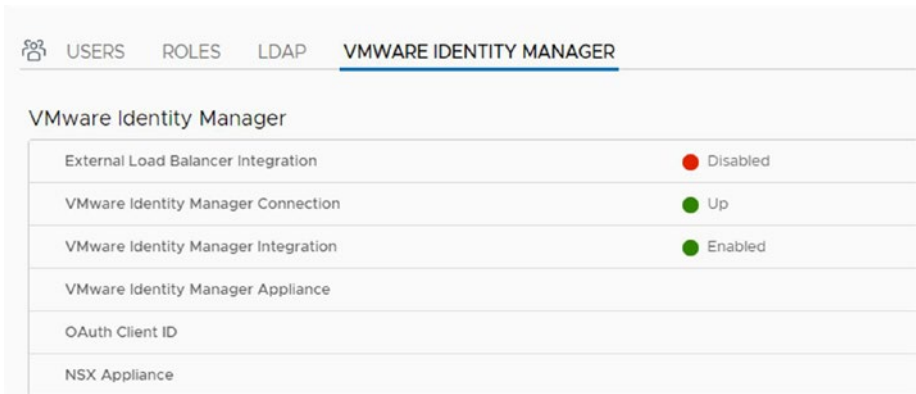


Figure 8-10. vIDM with NSX-T integration (NSX-T side) verification

Logging In with vIDM Services Enabled

When integration with vIDM is enabled, you are redirected to the VMware Identity Manager login page for authentication.

Default NSX-T GUI Login Screen

The default login page can also appear if integration with vIDM is configured, but vIDM is down or not reachable at the time of the login for whatever reason. See Figure 8-11.

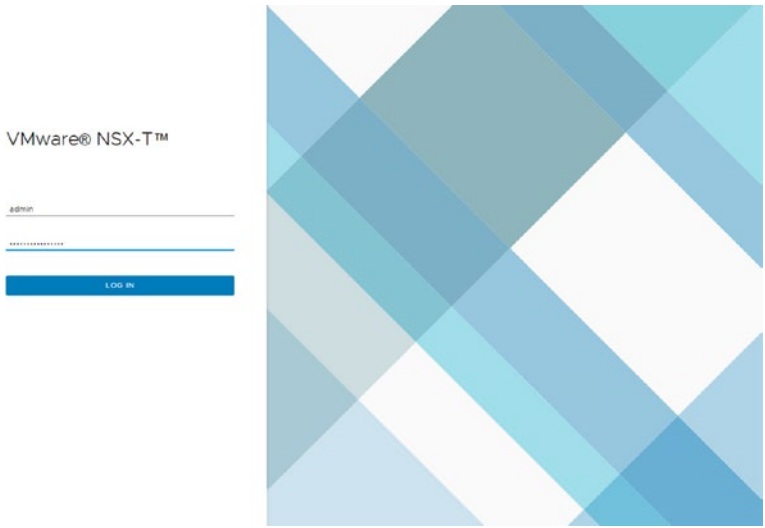


Figure 8-11. Default login page on NSX-T (vIDM not reachable)

Local Login When VIDM Is Enabled

You can log in using a local user, thereby bypassing the vIDM (troubleshooting or administration). See Figure 8-12. Browse to the following URL:

`https://<NSX_Manager_FQDN>/login.jsp?local=true.`

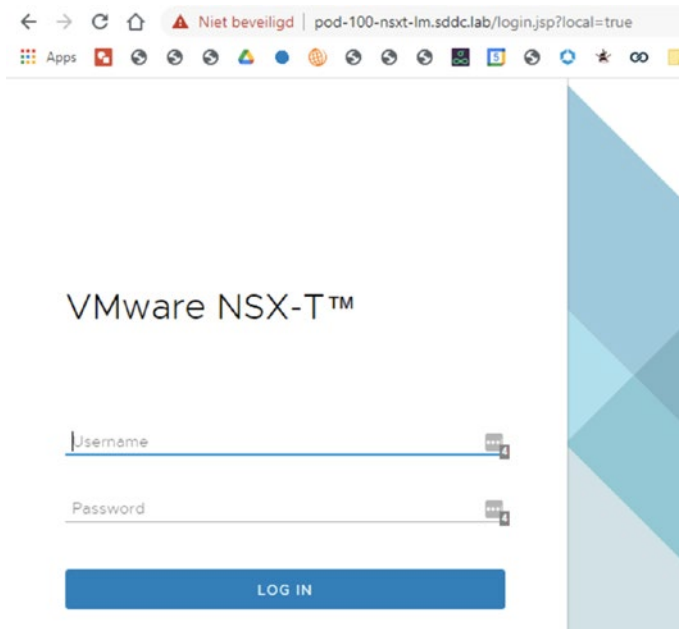


Figure 8-12. Default login page on NSX-T (vIDM reachable but local login URL is used)

NSX-T LDAP Integration

This section explains the integration between LDAP and NSX-T.

LDAP

LDAP is a protocol for accessing and managing distributed directory services. Distributed directory services are used to store information about users and groups, the network infrastructure, and network services. NSX-T supports Active Directory over LDAP and OpenLDAP.

LDAP and NSX-T Integration Benefits

Integrating LDAP with NSX-T makes it possible to use the existing directory service infrastructure. You don't have to deploy the vIDM appliance.

Besides these benefits, the integration is also simple to configure and easy to manage.

LDAP Authentication and NSX-T

When LDAP is integrated successfully with NSX-T, users are authenticated using the LDAP database's user information.

The authentication is shown in Figure 8-13 with these steps described:

1. A user initiates a login request from the browser.
2. The NSX-T Manager receives this login request and creates an LDAP bind request to the identity source (for example, Active Directory).
3. The identity source then returns an LDAP bind response to the NSX-T Manager.
4. Based on the LDAP bind response, the NSX-T Manager authenticates the user and displays the NSX-T home page or displays an authentication error.




Figure 8-13. LDAP authentication steps

Identity Source Configuration

You can add a new identity source by navigating to System ► Settings ► Users and Roles ► LDAP. See Figure 8-14.

When you configure a new LDAP identity source, you can specify the parameters outlined in Table 8-6.

Table 8-6. LDAP Identity Source Configuration Parameters

Parameter	Description
Name	A name for the identity source.
Domain Name	The domain you want to add as an identity source.
Type	Active Directory over LDAP and OpenLDAP are both supported identity source types to choose from.
LDAP Server	Specify the connection settings to your LDAP server.
Base DN	Defines the (starting) point from where a server searches for users.

The screenshot shows the VMware Identity Manager interface for configuring an LDAP identity source. The top navigation bar includes 'USERS', 'ROLES', 'LDAP', and 'VMWARE IDENTITY MANAGER'. Below the navigation, there is a button 'ADD IDENTITY SOURCE' and a 'COLLAPSE ALL' option. The main content area displays a table with the following columns: Name, Domain Name (FOQN), Type, LDAP Servers, and Connection Status. Below the table, a form is shown for configuring a new identity source. The form includes the following fields and values:

- Name:** IDAP sddc lab
- Domain Name (FOQN):** Example: vmware.com
- Type:** Active Directory over LDAP
- Base DN:** Example: CN=Users,DC=sddc,DC=lab
- Alternative Domain Names:** Enter Alternative Domain Names
- Description:** Example: vmware.com

At the bottom of the form, there are 'SAVE' and 'CANCEL' buttons.

Figure 8-14. LDAP identity source

LDAP Server Configuration

To configure the LDAP server, you need to specify the mandatory fields when configuring the identity source. When you click on Set, you can configure the LDAP server.

When you configure a new LDAP server, you can specify the parameters outlined in Table 8-7.

Table 8-7. *LDAP Server Configuration Parameters*

Parameter	Description
Hostname/IP	The fully qualified domain name or IP address of the LDAP server. For LDAPS configurations, the FQDN must match the hostname in the LDAP server certificate.
LDAP Protocol	Select LDAP (unsecured) or LDAPS (secured).
Port	The default LDAP port 389 and LDAPS port 636 are used for the directory sync. You should not change the default values.
Use StartTLS	This is available only for the LDAP protocol. SSL/TLS is used to establish a secure connection.
Certificate	If the LDAPS or the UseStartTLS options are configured, you can select the SHA-256 thumbprint that is suggested by NSX-T Manager or you can enter an SHA-256 thumbprint manually.
Bind identity	The (AD) domain account with read permission for all objects in the domain tree.
Password	The password for the (AD) bind identity account.

To verify that you can connect to the LDAP server, click Check Status, as shown in Figures 8-15 and 8-16.

Set LDAP Server ×

Identity Source #Ldap Servers 1

ADD LDAP SERVER Maximum: 1

Hostname/IP	LDAP Protocol	Port	Connection Status
ldap.sddc.lab * ⓘ	LDAP ▼	389 *	Check Status
Use StartTLS <input type="checkbox"/>		Certificate	Enter Certificate
Bind Identity Enter Bind Identity Format: user@domainName or specify the distinguished Name		Password 🔒
ADD CANCEL			

CANCEL APPLY

Figure 8-15. LDAP server configuration (1)

USERS ROLES **LDAP** VMWARE IDENTITY MANAGER ?

ADD IDENTITY SOURCE COLLAPSE ALL Filter by Name, Path and more

Maximum: 3

⋮	Name	Domain Name (FGDN)	Type	LDAP Servers	Connection Status
▼	IDENTITY-SOURCE	ldap.sddc.lab	Active Directory over LDAP	1	Check Status
	Base DN	CN=Users,DC=sddc,DC=lab	Alternative Domain Names		
	Description				

Figure 8-16. LDAP server configuration (2)

GUI Login Using LDAP

You log in as an Active Directory or as an OpenLDAP user by specifying the domain name on the NSX login page. See Figure 8-17.



Figure 8-17. NSX login page (Active Directory or OpenLDAP user)

Role-Based Access Control (RBAC)

This section discusses the role-based access control (RBAC) features offered by NSX-T.

NSX-T User Account Types

You can log into NSX-T using a local user, users managed by vIDM or Active Directory, and OpenLDAP users.

Note vIDM can also use Active Directory (LDAP) as an identity source.

Policy Management, Access, and Authentication

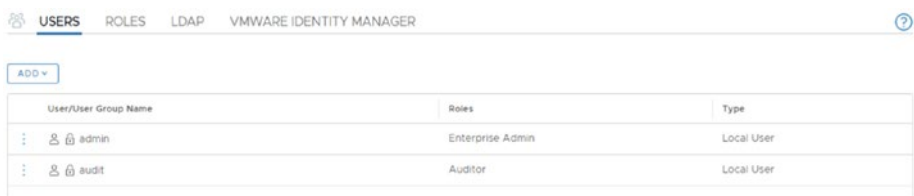
User access and the authentication policy management is different for local users and for external users. These differences are outlined in Table 8-8.

Table 8-8. *Authentication Policy Management*

Local users	This account type has access to the NSX-T Manager through all available interfaces: CLI, API, and UI. This authentication policy is configured directly from the NSX-T Manager CLI.
Users managed by vIDM	This account type can access the NSX-T Manager through the CLI and UI. The authentication policy is configured from the vIDM administration console. vIDM users can access/authorize the NSX-T API as well.
Active Directory and OpenLDAP users	This account type can access NSX-T Manager through the CLI and UI. This authentication policy is configured from the Active Directory or OpenLDAP. AD/OpenLDAP users can access/authorize the NSX-T API as well.

Local NSX-T Users

By default, NSX-T is preconfigured with two built-in users—`admin` and `audit` (Figure 8-18). The NSX-T Manager and NSX-T Edge transport nodes also have a preconfigured root user for CLI access.



The screenshot shows the 'USERS' tab in the NSX-T Manager interface. It displays a table with two rows of user information. The first row is for the 'admin' user, which has the role 'Enterprise Admin' and is a 'Local User'. The second row is for the 'audit' user, which has the role 'Auditor' and is also a 'Local User'. There is an 'ADD' button and a search icon at the top of the table.

User/User Group Name	Roles	Type
admin	Enterprise Admin	Local User
audit	Auditor	Local User

Figure 8-18. *Local NSX-T users*

Password Change

When you use the NSX-T CLI, you can change the password for (local) admin and audit users using the following command:

```
set user <username> [password <password> [old-password <old-  
password>]]
```

The password needs to meet the following requirements:

- At least 12 characters in length
- At least one uppercase character
- At least one lowercase character
- At least one numeric character
- At least one special character

Authentication Policy Settings Configuration

You can change the authentication policy settings for local users using the NSX CLI. The commands are listed in Table 8-9.

Table 8-9. *Authentication Policy Settings Change Commands*

Action	Commands
Set the minimum password length	set auth-policy minimum-password-length <password-length>
Set the GUI and API authentication policies	set auth-policy api lockout-period <lockout-period> set auth-policy api lockout-reset-period <lockout-reset-period> set auth-policy api max-auth-failures <auth-failures>
Set the CLI authentication policy	set auth-policy cli lockout-period <lockout-period> set auth-policy cli max-auth-failures <auth-failures>

Authentication Policy Setting Configuration for VIDM Users

You can configure the authentication policy settings for vIDM users. In the vIDM administration console, navigate to Identity & Access Management ► Policies. See Figure 8-19.




Figure 8-19. Authentication Policy Settings

Authentication Policy Setting Configuration for LDAP Users

You can use standard Windows or Linux tools to configure the authentication policy settings for Active Directory and OpenLDAP users. See Figure 8-20.

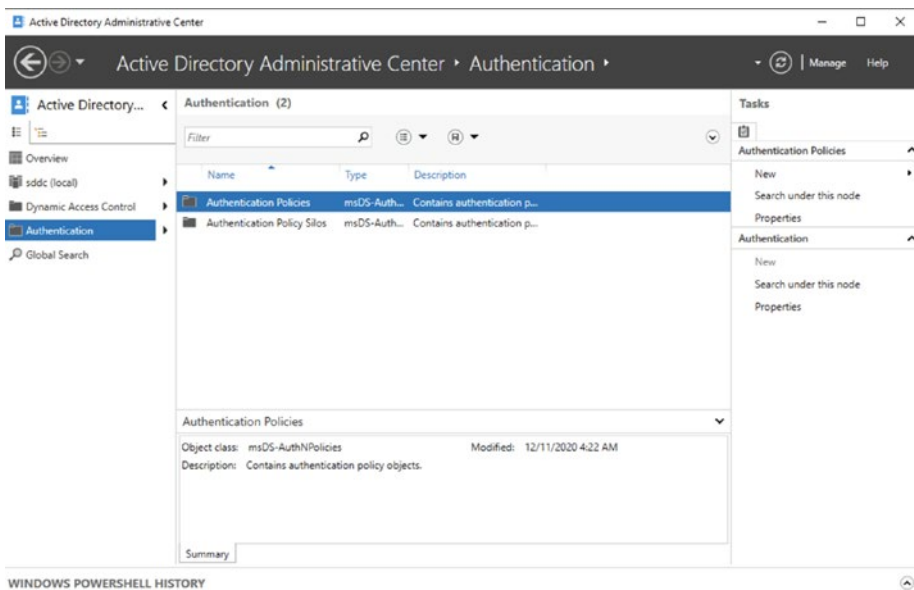


Figure 8-20. Windows tool to configure the authentication policy settings

Default Roles

NSX-T provides the built-in roles described in Table 8-10.

Table 8-10. *NSX-T Default Roles*

Role	Description
Auditor	Read permissions on all features
Enterprise Administrator	Full access permissions on all features
Load Balancer Administrator	Read permissions on all networking services and full access permissions on load-balancing features
Load Balancer Auditor	Read permissions on all networking services and load-balancing features
Network Engineer	Full access permissions on all networking services
Network Operator	Read permissions on all networking services and execute permissions on monitoring and troubleshooting tools
Security Engineer	Full access permissions on all security configurations
Security Operator	Read permissions on all security configurations and execute permissions on monitoring and troubleshooting tools
GI Partner Administrator	Role used for third-party endpoint protection service insertion
Netx Partner Administrator	Role used for third-party network Introspection service insertion
VPN Administrator	Read permissions on all networking services and full access permissions on VPN features

Role Assignment for Local Users

The local users are preconfigured with specific roles that cannot be modified. The `admin` user is preconfigured with the Enterprise Administrator role. The `audit` user is preconfigured with the Auditor role.

Figure 8-21 shows an example with the `audit` user, where I cannot add an additional Tier-1 gateway.




Figure 8-21. Audit user role assigned

Role Assignment for vIDM Users

You can add, change, or delete role assignments for vIDM users or user groups with the steps outlined in Table 8-11. See Figure 8-22.

Table 8-11. Role Assignment Configuration Steps (vIDM)

Step	Description
1	In the NSX-T GUI, select System ► Settings ► Users and Roles ► Users.
2	Click Add and select Role Assignment for vIDM.
3	Search for the users or user groups that you want to assign the roles to.
4	Select a role or roles and click Save.



Figure 8-22. Role assignment to user (vIDM)

Role Assignment for LDAP Users

You can add, change, and delete role assignments for LDAP users or user groups with the steps outlined in Table 8-12. See Figure 8-23.

Table 8-12. Role Assignment Configuration Steps (LDAP)

Step	Description
1	Select System ► Settings ► Users and Roles ► Users.
2	Click Add and select Role Assignment for LDAP.
3	Select the identity source.
4	Search for the users or user groups that you want to assign the roles to.
5	Select a role or roles and click Save.




Figure 8-23. Role assignment to user (LDAP)

Summary

In this chapter, you learned the purpose of the VMware Identity Manager (vIDM), LDAP, and integration with NSX-T. You also learned about NSX-T's local accounts and how to map different roles to different sources (vIDM and LDAP) in order to assign rights based on policies.

The next chapter explains NSX-T Federation.

CHAPTER 9

NSX-T Federation

This chapter explains how to use NSX-T stretched networking and stretched security features across multiple sites via NSX-T Federation.

NSX-T Federation

NSX-T Federation provides multisite network and security functionality for different data center locations. The number of sites that are supported may be different based on the NSX-T release and can be checked from the website <https://configmax.vmware.com/>. For NSX-T 3.0, this is three sites and for NSX-T 3.1, this is four sites.

The multiple sites can be adopted for several use cases:

- Providing high availability for your applications
- Providing a better application response time
- Providing the most cost-effective hosting solution depending on how critical the applications are
- Providing a unified network and security configuration during mergers or acquisitions

NSX-T Federation Use Cases

NSX-T Federation helps deliver a cloud-like operating model by simplifying the consumption of networking and security constructs.

With NSX-T Federation, you can simplify your operations and management because you have a consistent network and security policy configuration.

Policy Consistency and Simplification of Operations

NSX-T Federation introduces new components and objects. The NSX-T Manager node that we are used to is now called a *Local Manager* (LM) and a new NSX-T Manager type is introduced, called a *Global Manager* (GM).

The GM is a centralized console that enables you to view and configure the Federation services from a single management user interface.

NSX-T Federation provides operational simplicity and consistent policy configuration. You can manage the network as a single entity while keeping configuration and operational state synchronized across multiple locations.

Security policies attach and move with the workload, ensuring that policy compliance is maintained during workload failover or migration between locations.

The stretched objects (such as a stretched Tier-0 or Tier-1 gateway) currently do not support all the features that (local) Tier-0 or Tier-1 gateways support. This may change in future releases.

The services that are not supported on a stretched (Federated) Tier-0 or Tier-1 gateway include:

- ID-FW, L7-FW, and IDS
- Load balancing
- VPN
- L2-bridge
- Service insertion/guest introspection
- VRF Lite

NSX-T Federation Components

This section elaborates on the NSX-T Federation components and discusses how these components interact.

Global Manager (GM)

The GM provides the GUI and the REST APIs to configure (stretched) objects across multiple geographical locations.

A single standalone GM node spans multiple locations, but it's primarily hosted on one site, with capabilities to protect against failures.

The GM offers the ability to perform global configurations that can span multiple locations or a single location. Any global configuration is pushed to the relevant LMs to complete the configuration of the stretched objects. Configuration that needs protection against location failures must be configured in a GM.

Local Manager (LM)

The LM manages a single location (not stretched). The GUI and API access are available by using the LM virtual IP address.

The LM realizes any global configuration that is sent by the GM. The LM will still accept user configuration that is required for the local location with the local objects. The LM owns the infrastructure preparation and configuration of the host and edge transport nodes, transport zones, IP pools, etc.

GM vs LM Components

The different capabilities of the GM and LM are listed in Table 9-1. See Figure 9-1.

Table 9-1. GM vs LM Capabilities

	GM	LM
Deployed in the VM form factor	X	X
Include a global control plane	-	-
Deployed as a cluster in each location for high availability and scalability	-	X
Override selective global objects that GM creates	-	X
Deployed as a standby cluster in another location for high availability.	X	-

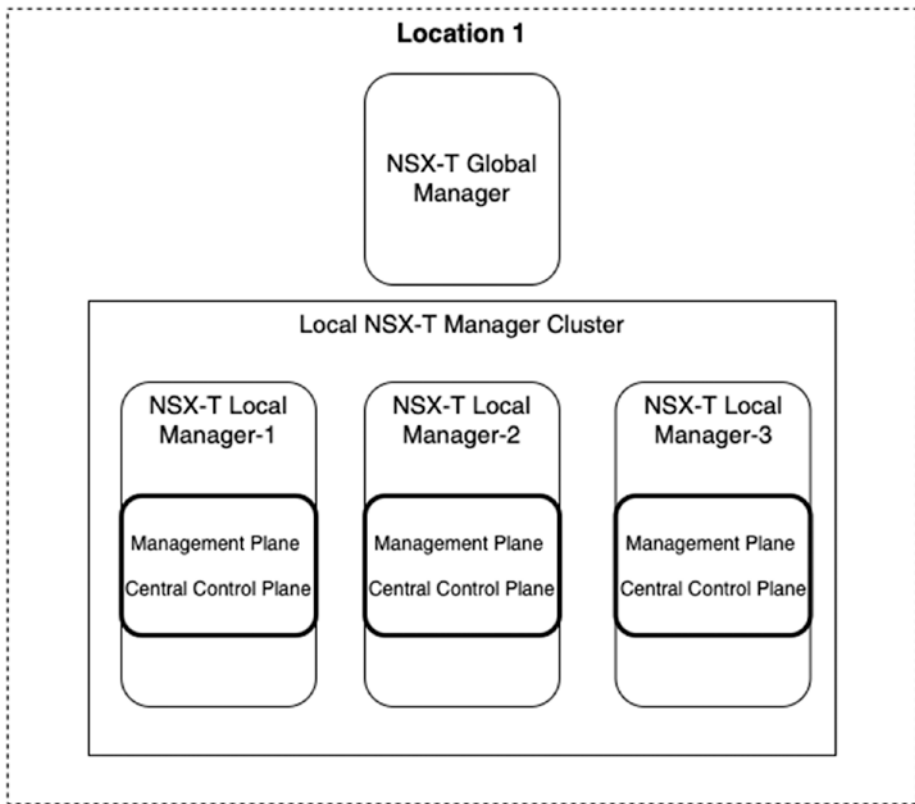


Figure 9-1. Global and Local Manager positioning

Multi-Site Federation View

Figure 9-2 shows three locations with a GM that are deployed in a cluster. That cluster consists of standalone VMs deployed in a single location with a standby cluster on another site in case of any failures. Each location has an LM cluster deployed. Each location is also managed independently with a dedicated vCenter server; the vSphere clusters are local to the site and are not stretched.

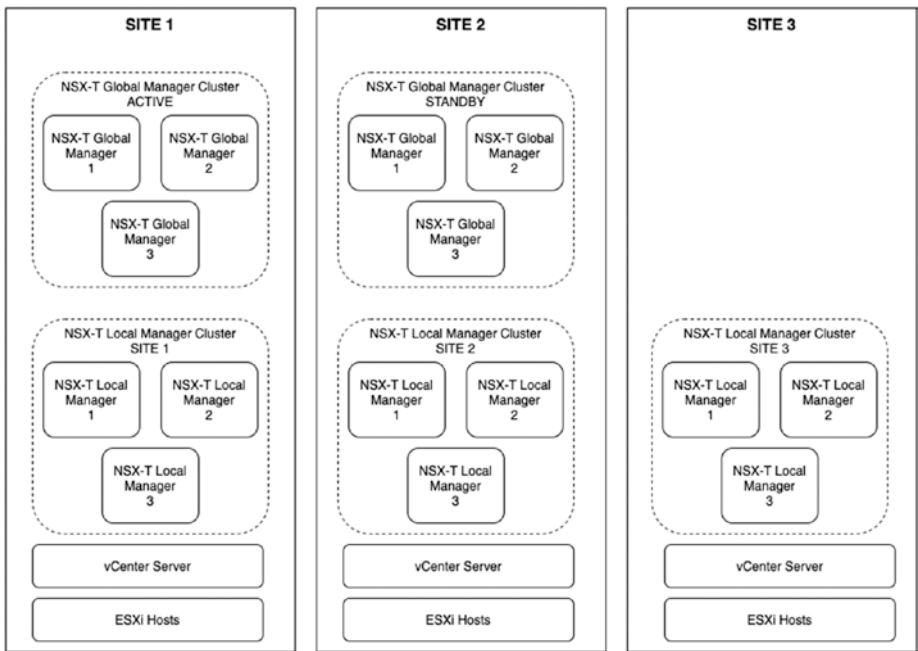


Figure 9-2. GM clustering option

Another way of deploying the GMs is to spread them across three locations, whereby each location hosts one GM instance. In order to do this, the latency between the three sites needs to be below 10 milliseconds. See Figure 9-3.

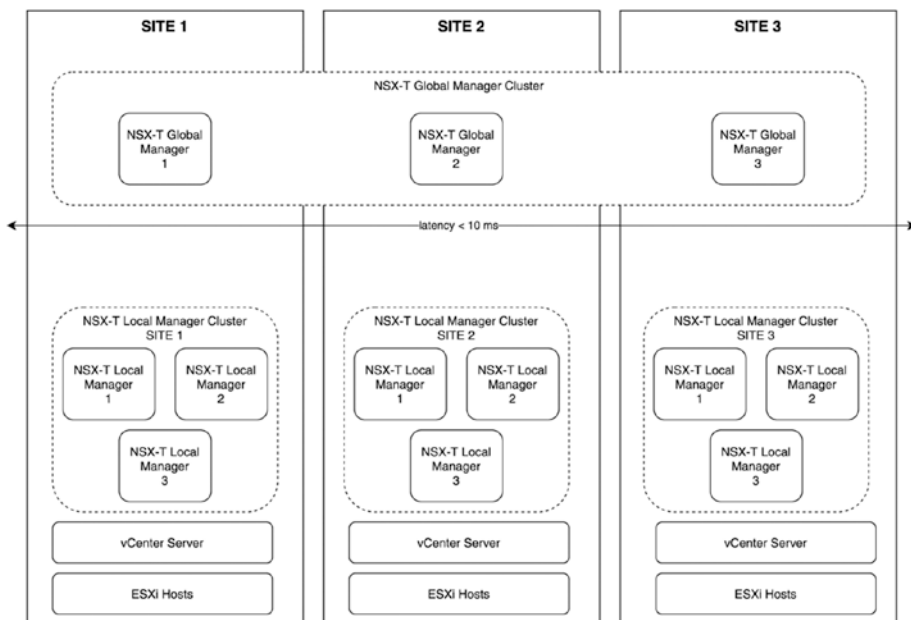


Figure 9-3. GM site local option

NSX-T Federation Consumption Methods

With NSX-T, you can configure both the GM and the LM. The Global Manager cluster receives the “global configuration” and the Local Manager cluster receives the “local configuration.” When a configuration is made in the Global Manager that involves federated objects, that configuration is pushed to the Local Manager cluster as well. See Figure 9-4.

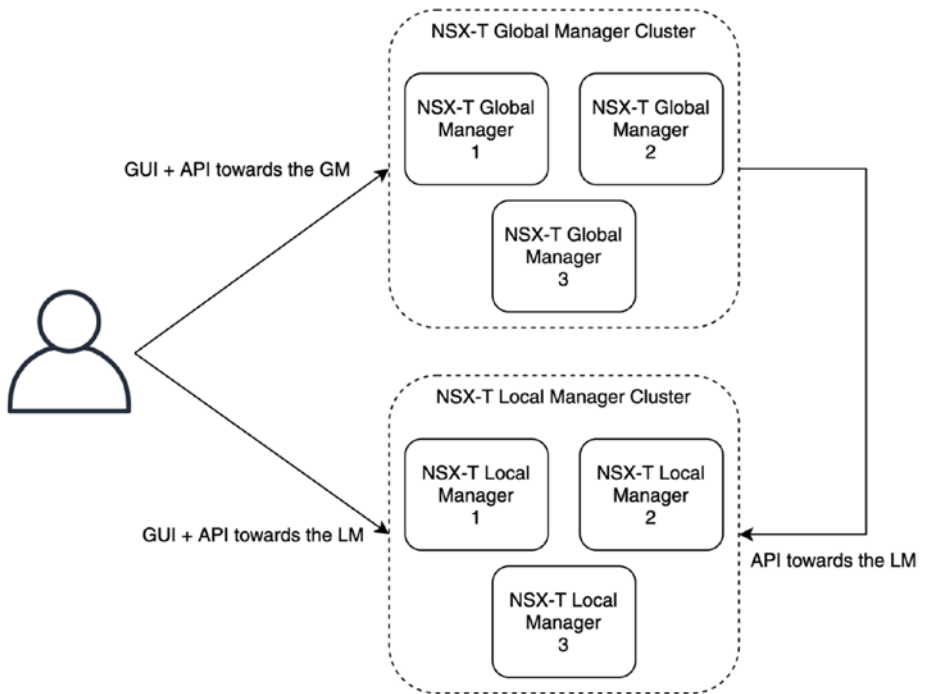


Figure 9-4. GM and LM consumption methods

LM Configuration Ownership

The LM owns the network objects that it creates, such as Tier-0 and Tier-1 gateways, segments, segment profiles, etc. All these (local) objects can be configured from the LM. Only the LM can modify and delete these (local) objects. Objects are not visible to the GM.

The LM always manages the Fabric, such as the host transport nodes, edge transport nodes, transport zones, etc. The LM completely owns the preparation of the infrastructure. Only the LM can create, modify, or delete these infrastructure components or objects.

GM Configuration Ownership

The GM owns the network objects that it creates, such as stretched Tier-0 and Tier-1 gateways, stretched segments, etc. All these (global objects) can be configured from the GM. Only the GM can modify or delete these (global) objects. The objects created by the GM are visible to the LM, but are in read-only mode.

LM Configuration

The configuration of the Local Manager (LM) is described with the following steps (Figure 9-5):

1. A user can send the configuration to each location LM as required.
2. Each LM (cluster) stores the configuration locally.
3. There is no configuration sent to the GM.

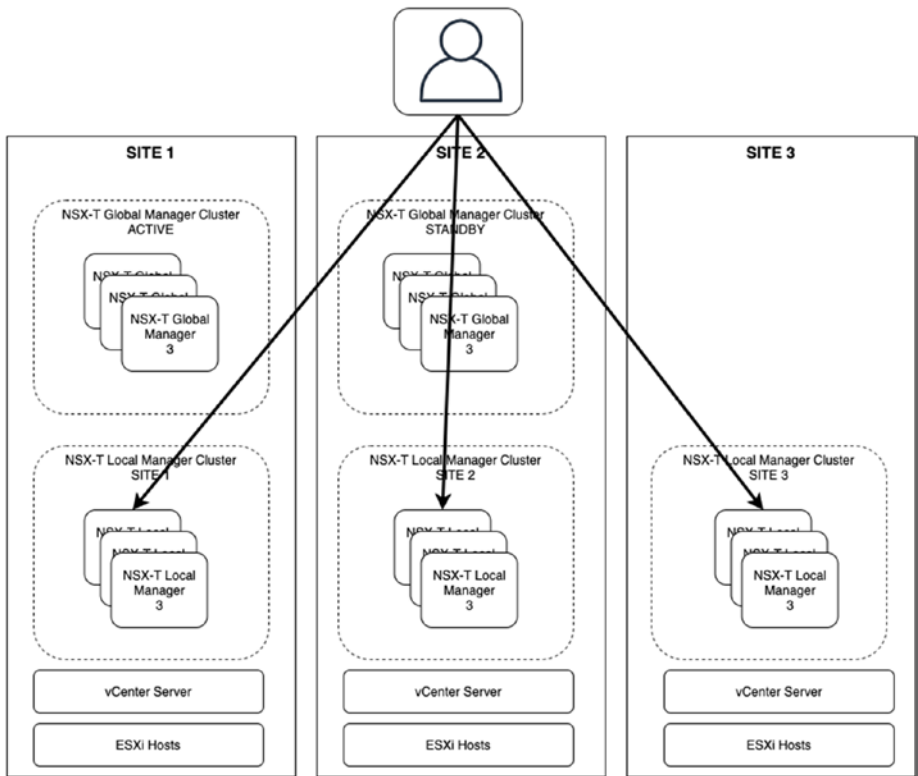


Figure 9-5. LM configuration flow

GM (Federation) Configuration

The configuration of the Global Manager (GM) is described in the following steps (Figure 9-6):

1. The user sends the configuration to the GM.
2. The GM sends it to the relevant LMs.

Note Some of the sites may not participate in the NSX-T Federation configuration, which means that the LM local to this site will not receive any configuration details from the GM.

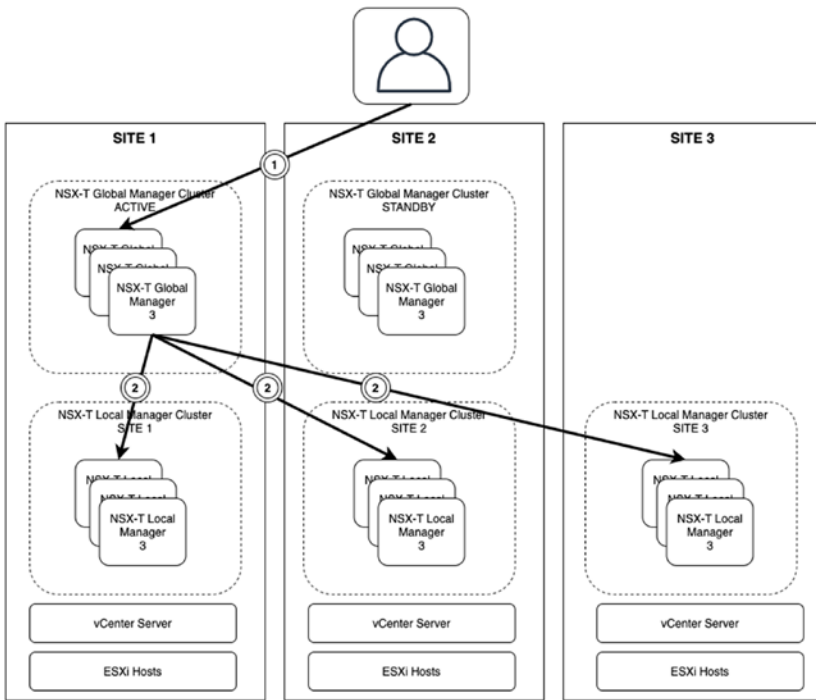


Figure 9-6. GM configuration flow

GM (Federation) Initial Configuration Steps

In the configuration example, networks are created and realized only in two locations (Site 1 and Site 2) (Figure 9-7):

- The Tier-0 and Tier-1 gateways are stretched across Site 1 and Site 2.
- The Blue-Segment associated with Tier-1 is also stretched to Site 1 and Site 2.
- Two virtual machines (Virtual Machine 1 and Virtual Machine 2) are connected to the Blue-Segment.
- This configuration is created on the GM and the LMs across (two) locations.

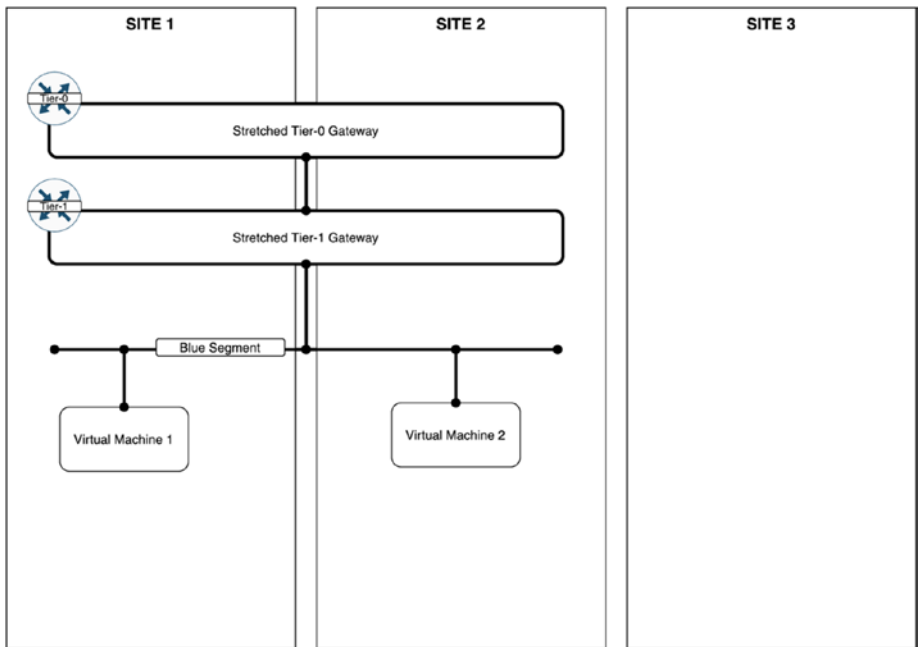


Figure 9-7. Stretched network across two sites

1. When submitting any configuration to the GM cluster in Site 1, this GM stores the configuration locally.
2. The GM then pushes the configuration to the LM clusters of Site 1 and Site 2 because the configuration is intended *only* for them.
3. The configuration is not sent to LM of Location 3, because the (stretched) objects are not associated with the configuration in Location 3. See Figure 9-8.

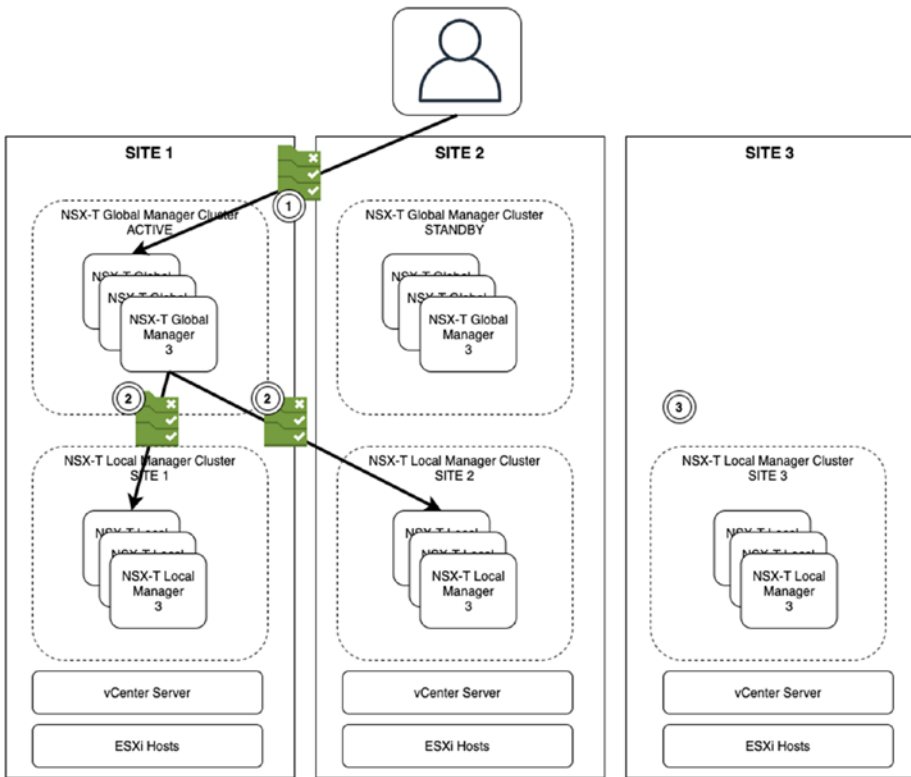


Figure 9-8. Configuration steps (1)

1. The LMs of each location learn the control plane information of the intended configuration.
 - a. The LM does not update the control plane information to the GM, because the GM stores only configuration.
 - b. The GM GUI enables you to see the inventory of all locations by querying the relevant LMs.
 - c. In this example, the control plane of the LM in each location learns the IP and MAC information of the virtual machines associated with the Blue-Segment.

2. The LMs of each location initiate the sync to exchange the control plane configuration of the topology.
3. Each LM stores the information of the topology discovered through the sync operation. The control plane of the LM in each location has the IP and MAC information of all the virtual machines associated with Blue-Segment across Site 1 and Site 2. See Figure 9-9.

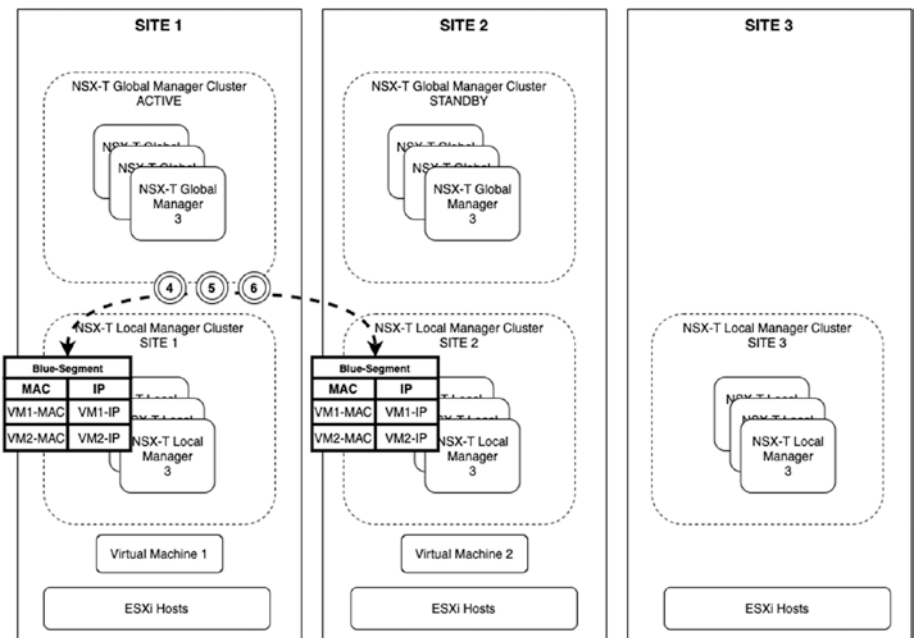


Figure 9-9. Configuration steps (2)

GM (Federation) Greenfield

This section discusses how to deploy NSX-T Federation in a greenfield environment.

Prerequisites

Before you can deploy NSX-T Federation, you need to follow the prerequisites listed in Table 9-2.

Table 9-2. *NSX-T Federation Prerequisites*

Site-to-site traffic	The latency (RTT) should be less than 150 ms between two sites.
IP and firewall connectivity	For the management traffic between the GM and the LM. For the data plane traffic between edge transport nodes (RTEP). No NAT on Remote Tunnel Endpoints (RTEPs).
WAN bandwidth requirement	No congestion for the management plane (GM-LM traffic). No congestion for the data plane.
WAN MTU requirement	An MTU of 1700 bytes or more to avoid the edge transport node RTEP traffic fragmentation.

Onboarding

You need to connect an LM to a GM to enable the Federation functionality. This is done using the following steps:

1. Establish connectivity and mutual authentication between the GM and the LM.
2. Optionally, you can configure RTEPs to permit the stretching of networks between sites.

The onboarding process is repeated a few times for each site until all the sites are registered. The onboarding process does not affect existing workloads operated by the LMs.

The First Location

For the first location (Site 1), you need to perform the following tasks:

- Deploy the LM and finish the required configuration.
- Deploy the GM.
- Complete the initial wizard to prepare the location for Federation:
 - The LM is linked to the GM.
 - A name is selected for the first location and the corresponding location object is created.
 - An (local) edge cluster, in charge of traffic across locations, is selected or deployed.
 - Optionally prepare the RTEP interfaces for all cluster members.

Additional Locations

For the remaining subsequent locations (Location n), complete these tasks:

- All the subsequent LMs can join the GM.
- A name is selected for each location and the corresponding location object is created.
- A (local) edge cluster is selected or deployed for each location to enable communication across locations.
- Optionally prepare the RTEP interfaces for all cluster members.

GM (Federation) GUI: Local Management Cluster

When you browse to the GM, the Location Manager tab provides you with information about the GM and LMs. The Location Manager view provides the GM node information and the Local Management cluster information across locations. See Figures 9-10 through 9-12.

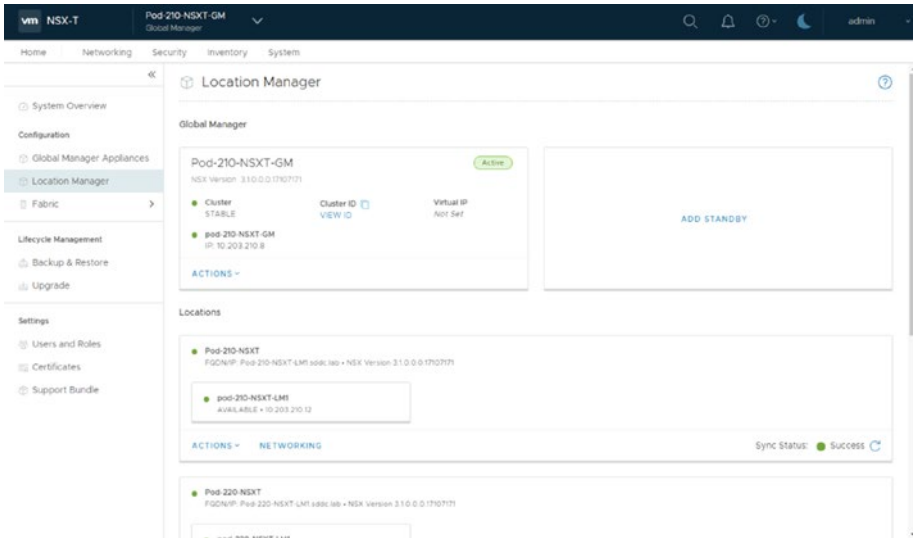


Figure 9-10. (GM) Location Manager view (1)

GM (Federation) GUI: Location Manager

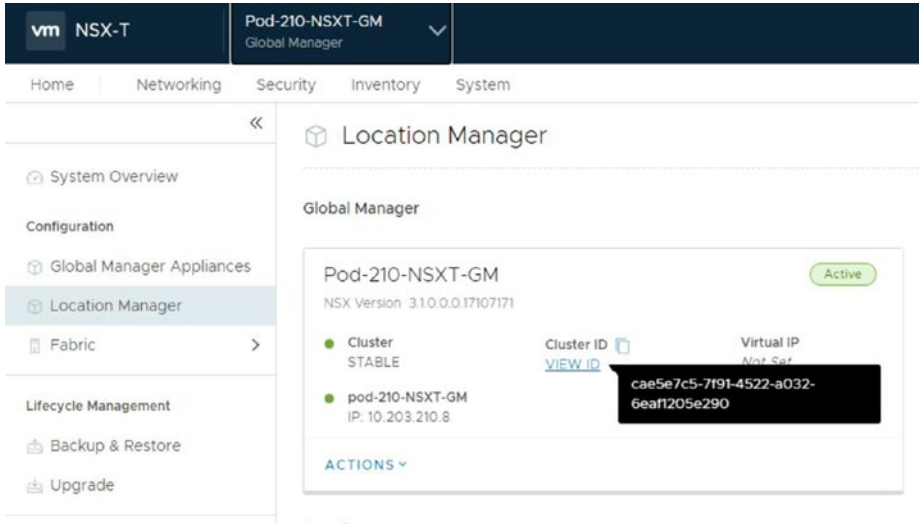


Figure 9-11. (GM) Location Manager view (2)

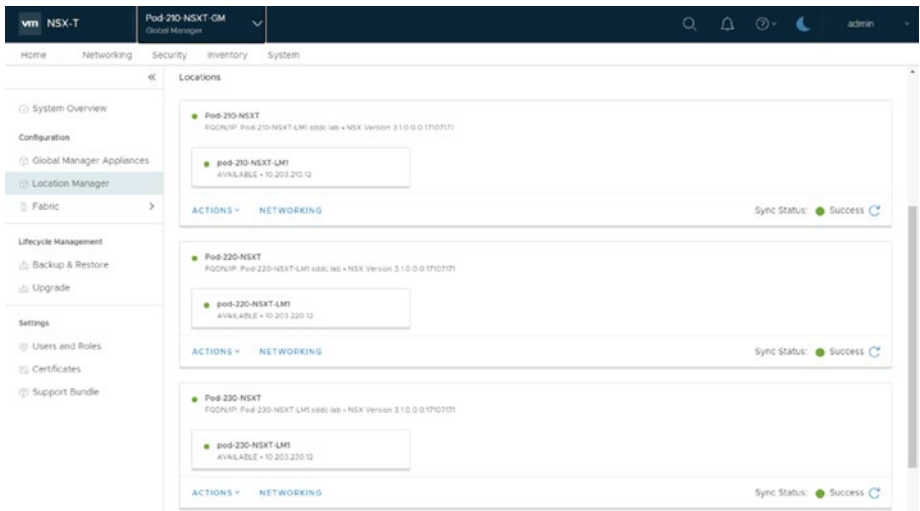


Figure 9-12. (GM) Location Manager view (3)

GM (Federation) GUI: Location Selector

Use the Location Selector menu to select the GM or different LMs, as shown in Figure 9-13.

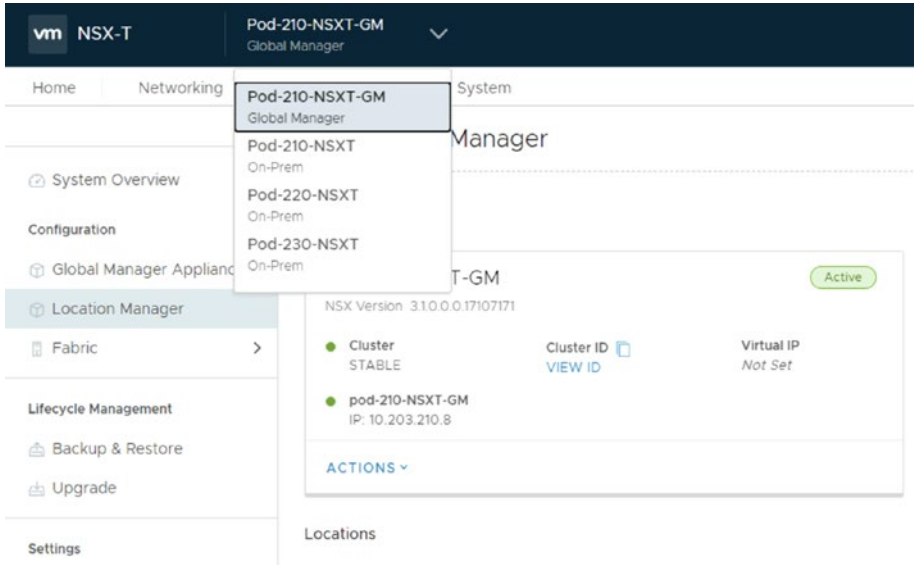


Figure 9-13. (GM) Location selector

GM (Federation) GUI: System Overview

The System Overview tab provides you with a complete overview of the Federation onboarded locations, as shown in Figure 9-14.

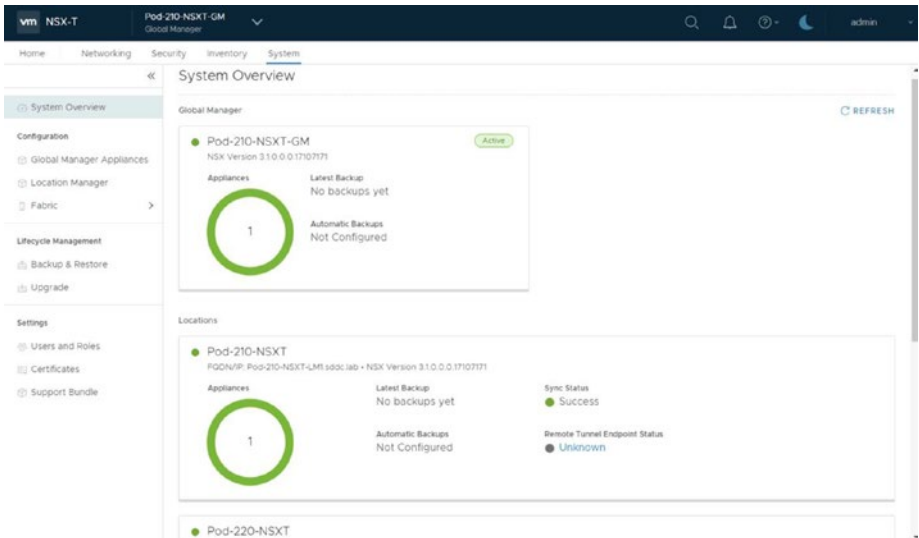


Figure 9-14. (GM) system overview

NSX-T Federation Networking

This section describes the stretched networking concepts in NSX-T Federation and explains the supported Tier-0 and Tier-1 gateway stretched topologies. It also describes the Tier-0 and Tier-1 gateway deployment modes in a multi-location scenario and explains the Layer-2 and Layer-3 packet walks in some deployment models.

NSX-T Federation Stretched Networking

NSX-T Federation supports stretching the Tier-0 and Tier-1 gateways and segment networking objects across multiple locations.

The services that are supported on a stretched Tier-1 gateway are NAT, gateway firewall, IPv6, DHCP, and DNS. See Figure 9-15.

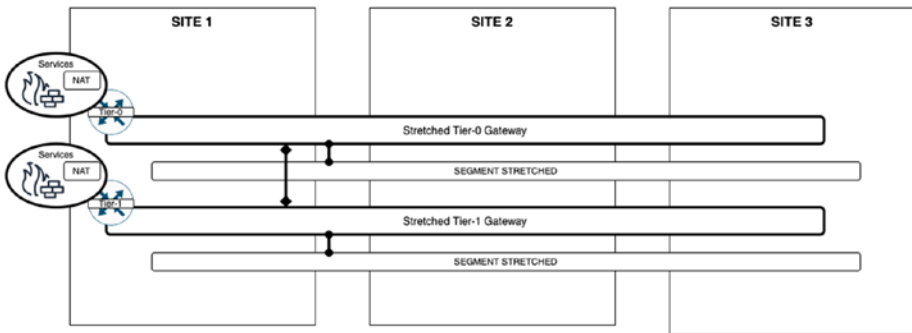


Figure 9-15. *Stretched networking*

The GM can stretch the Tier-0 and Tier-1 gateways across multiple locations, where all segments that are connected by using the downlink port are automatically stretched.

The edge transport nodes are used to forward cross-location traffic. This method eliminates the requirement for tunnels between the hypervisors across the different locations. The edge transport nodes that provide cross-location communication use RTEP interfaces.

LMs create Remote Tunnel Endpoints (RTEPs) on the edge transport nodes to route the traffic across sites. NSX-T does not encrypt the traffic, but you can encrypt this cross-location traffic (on the physical network level). See Figure 9-16.

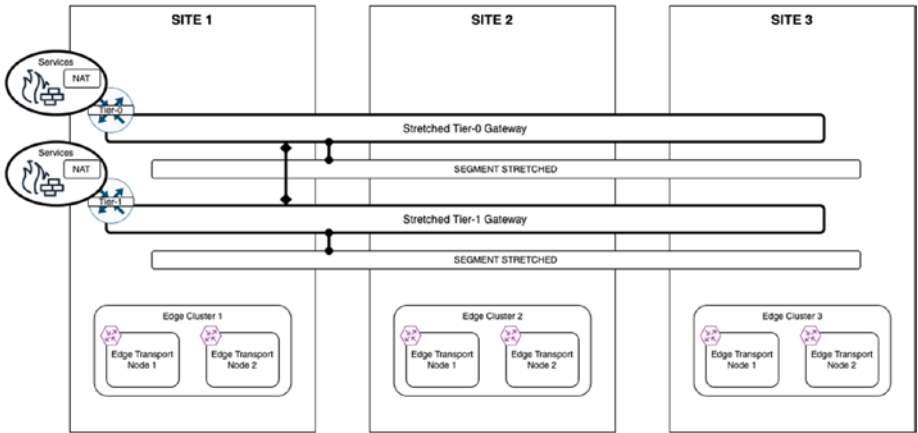


Figure 9-16. Remote tunnel endpoints (on edge transport nodes)

Logical Topologies with Tier-0 and Tier-1 Gateways

The GM supports stretch networks for cross-location communication and non-stretch networks for a local location.

The GM supports Tier-0 and Tier-1 gateways that are *not* stretched and that are local to a location. Segments that are associated with the Tier-0 and Tier-1 gateways are also *not* stretched and are local to a location. The span of a segment is equal to the span of the Tier-0 and Tier-1 gateways. See Figure 9-17.




Figure 9-17. Non-stretched networks

The GM also supports Tier-0 and Tier-1 gateways that can be stretched to all or a subset of the locations. The segments associated with stretched Tier-0 and Tier-1 gateways are also stretched to the same span. The span of a segment is equal to the span of Tier-0 and Tier-1 gateway.

If you run a Tier-1 gateway without any services (also called DR-only), the span of T1 is equal to the span of T0. See Figure 9-18.

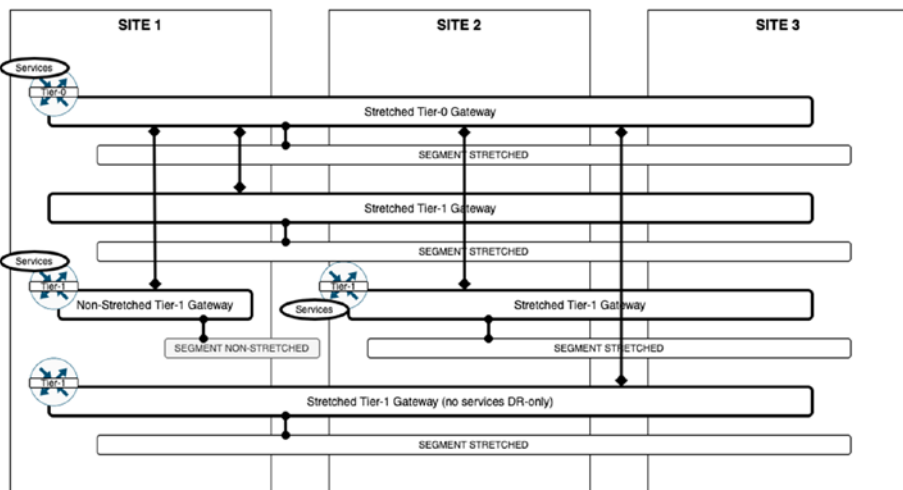


Figure 9-18. Stretched networks

A stretched Tier-0 gateway can connect to a non-stretched Tier-1 gateway, if the scope of the Tier-1 gateway is equal to or a subset of the Tier-0 gateway.

In Figure 9-19, the spans of the Tier-0-stretched gateway and the Tier-1-stretched gateway are not the same. This is not possible. The span of the Tier-1-non-stretched is a subset of the Tier-0-stretched span. This connection is possible.

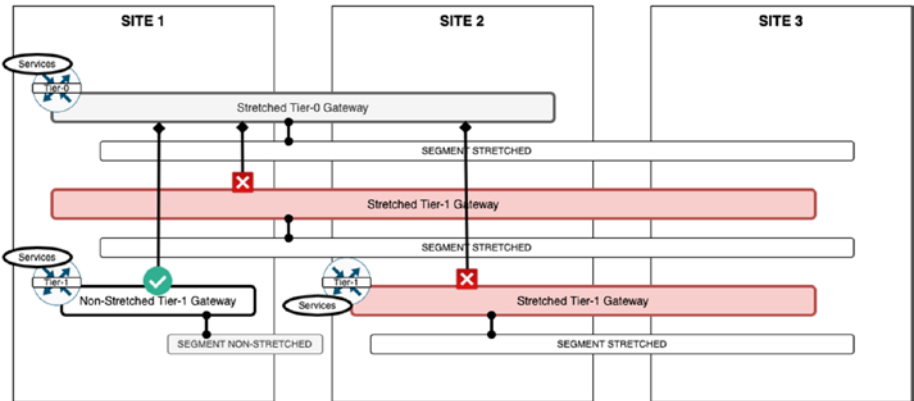


Figure 9-19. Span capabilities of stretched networks

Tier-0 and Tier-1 Gateway Deployment Rules

A Tier-0 gateway can be deployed in Active/Active (A/A) mode in each location. A Tier-1 gateway can only be deployed in Active/Standby (A/S) mode in each location.

A Tier-0 and Tier-1 gateway can be stretched across multiple locations. Each location can be configured in PRIMARY and SECONDARY for a Tier-0 and Tier-1 gateway.

The edge transport nodes (in an edge cluster) in each location back up the Tier-0 and Tier-1 gateways.

You must configure RTEPs on the edge transport nodes if you want the workloads to communicate across the different locations.

For the segments that are connected to a Tier-0 or Tier-1 gateway, the edge cluster must be available for the stretching segments.

Tier-0 and Tier-1 Single Location Deployments

The GM supports the Tier-1 gateway that is deployed in Active/Standby mode in a location. One edge transport node is active, and one edge transport node is in standby. The traffic ingresses or egresses from the active edge transport node of the cluster (and not from the standby edge transport node).

The GM also supports a Tier-0 gateway that is deployed in Active/Active mode in a location. All edge transport nodes are active. The traffic will ingress or egress from all the active edge transport nodes of the cluster. See Figure 9-20.

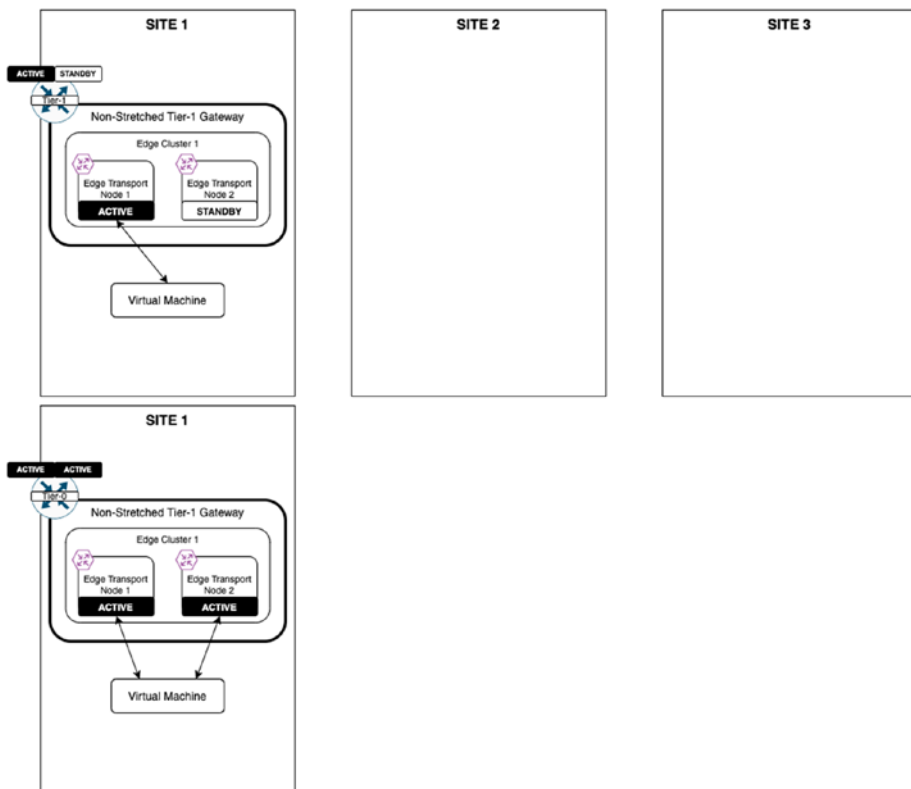


Figure 9-20. Single location Active/Standby

Tier-0 and Tier-1 Multi-Location Deployments

A location is defined and configured as either *primary* or *secondary*.

When you are using the primary and secondary deployment model, only one location can be configured as primary and all other locations are secondary.

Tier-0 and Tier-1 gateways can be deployed in primary and secondary (P/S) locations. The traffic moves to a Tier-0 gateway from a Tier-1 gateway and segments are contained inside the location. Only one location is active to send northbound egress traffic for a destination toward the physical network. This is illustrated in Figure 9-21, where Site 1 is used as the primary location and Sites 2 and 3 are the secondary locations.

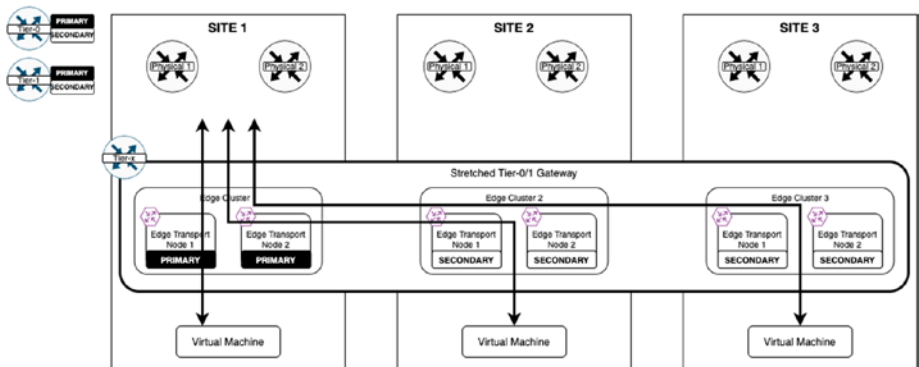


Figure 9-21. Multi-location primary/secondary

The Tier-0 gateway can be deployed in all primary locations mode. Traffic to the Tier-0 gateway from the Tier-1 gateway and segments is contained in the location. All locations are active to send northbound egress traffic toward the physical network. This is illustrated in Figure 9-22, where Sites 1, 2, and 3 are all used as primary locations.

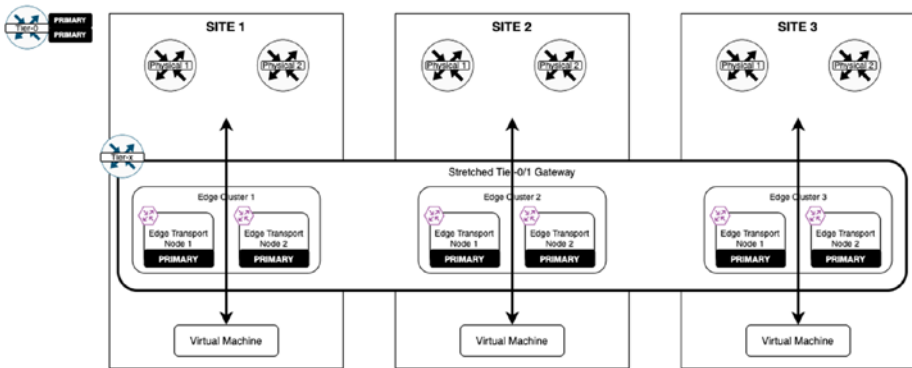


Figure 9-22. Multi-location all-primary

Tier-0 Multi-Location Stretched Modes

In Figure 9-23, the Tier-0 gateway is deployed in Active/Active mode in primary/secondary locations. For each location, two or more edge transport nodes can be in active mode. Active edge transport nodes in both primary and secondary locations can receive southbound traffic for their respective locations. Only the primary locations that have active edge transport nodes can send the northbound traffic of both locations.

The Tier-0 gateway does not support services in this deployment mode, but stateless NAT can be used.

Site 1 is primary, and Site 2 is secondary, where the stretched Tier-0 gateway is deployed in Active/Active mode for both locations. The RTEP IPs are configured on the edge transport nodes of all locations. Edge Transport Nodes 1 and 2 at Site 2 egress northbound traffic to either Edge Transport Node 1 or Edge Transport Node 2, or both on Site 1.

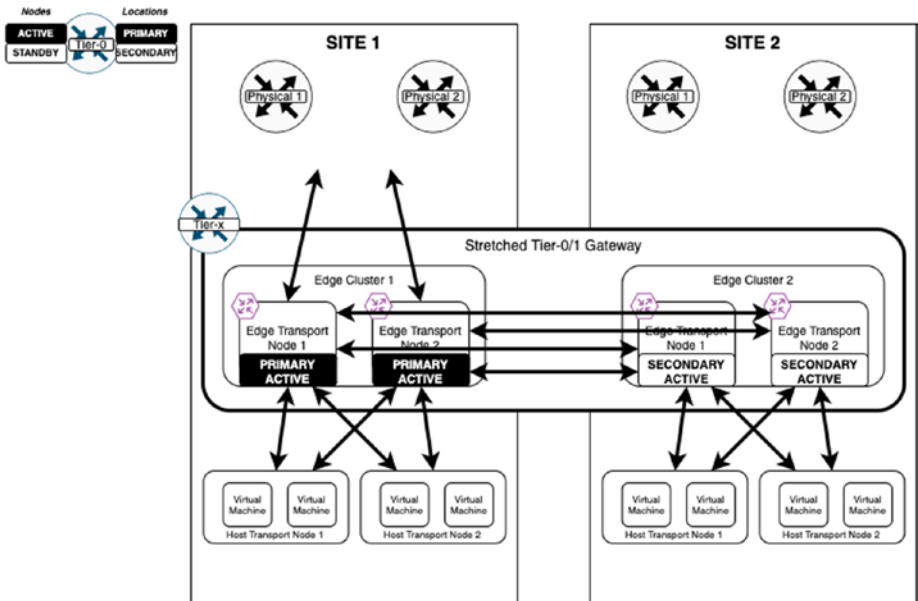


Figure 9-23. Multi-location primary/secondary (sites) plus active/active edges)

In Figure 9-24, the Tier-0 gateway is deployed in Active/Active mode in “all primary mode” across the locations. For each location, two or more edge transport nodes can be in active mode. Active edge transport nodes in both primary and secondary locations can receive both northbound and southbound traffic for their respective locations. This configuration is also called Active/Active Local Egress. The Tier-0 gateway does not support services in this deployment mode, but stateless NAT can be used.

In Figure 9-24, both locations are primary where the Tier-0 gateway is deployed in Active/Active mode for both locations. In this use case for local egress, all Tier-0 gateways route the northbound traffic locally to the local network.

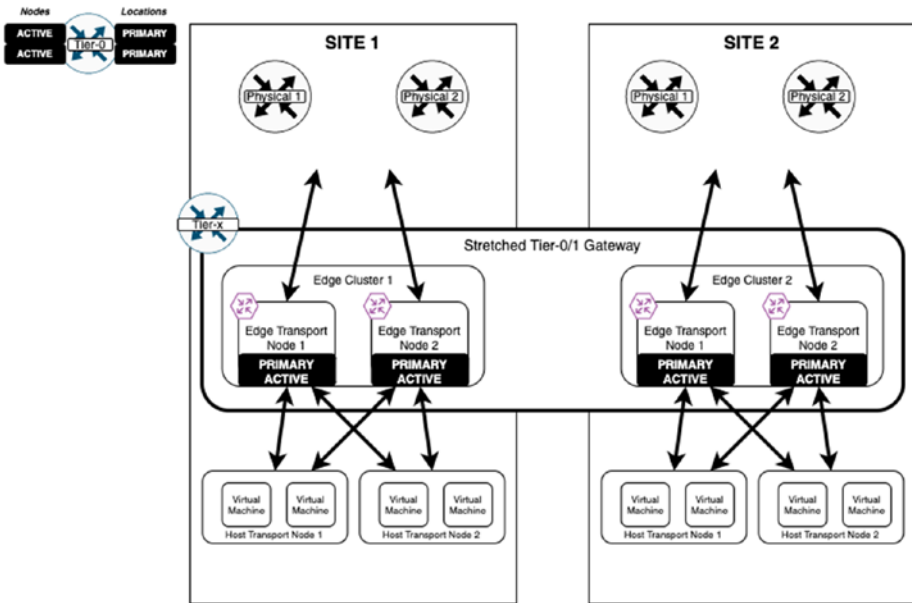


Figure 9-24. Multi location all-primary (sites) + active/active edges

Tier-1 Multi-Location Stretched Modes

In Figure 9-25, the Tier-1 gateway is deployed without any services. This is also called DR-only-mode. The Tier-1 gateway is deployed without services in all locations. Only the Distributed Router (DR) is created; the service router (SR) does not exist. Tier-1 gateways will route the northbound traffic directly to T0-stretched gateway on the hypervisor.

Workloads are connected to the T1-stretched gateway across locations. The communication happens through the edge transport nodes (using the RTEP interfaces) configured for L2 Stretch.

The T1-no-services gateway does not require edge transport nodes for routing, but the edge is required for L2 stretching. See Figure 9-25.

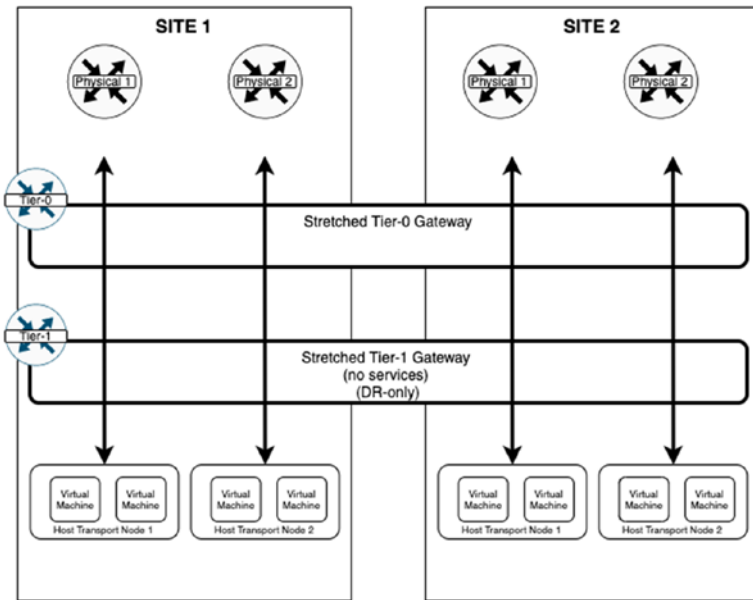


Figure 9-25. *Stretched Tier-1 DR-only (no services)*

In Figure 9-26, the Tier-1 gateway is deployed in Active/Standby mode in the primary/secondary mode across locations. Services are enabled on the Tier-1 gateway in this deployment mode. The active edge transport nodes in both primary and secondary locations can receive southbound traffic for their respective locations.

Only the primary location's active edge transport node can send the northbound traffic of both locations to the stretched Tier-0 gateway. See Figure 9-26.

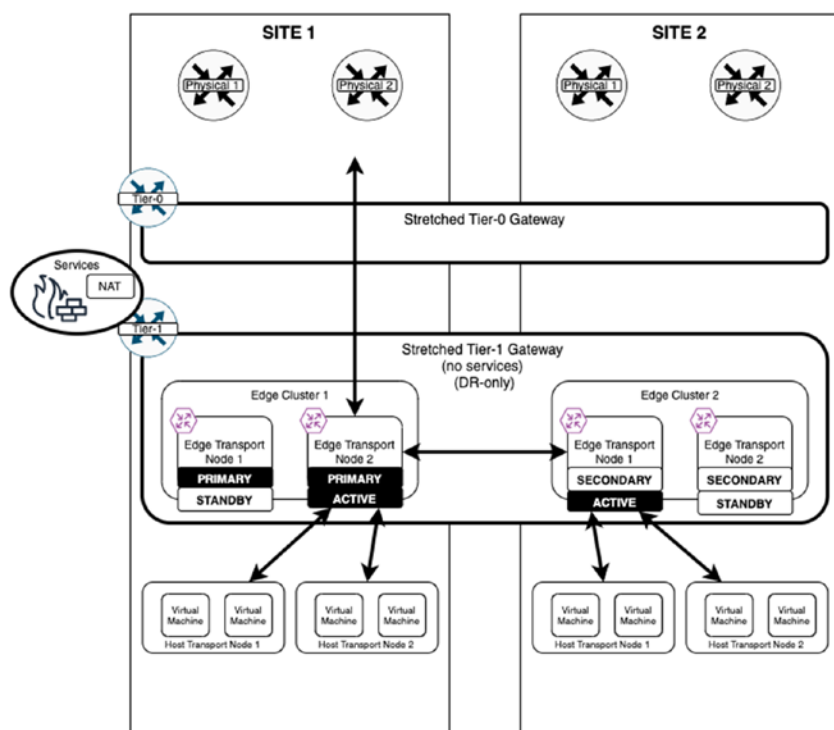


Figure 9-26. *Stretched Tier-1 (with services)*

Remote TEP (RTEP)

I have mentioned the term RTEP a few times now, but what is an RTEP exactly?

An RTEP is the edge transport node IP address that is used for remote edge transport node communication. The RTEPs use Geneve encapsulation for communication. Fragmentation on RTEP interfaces is allowed. However, for best performance, do not use fragmentation and try to configure an MTU of 1700 bytes between the full path of the edge transport nodes.

Stretched Layer-2 Network

Cross-location Layer-2 (broadcast domain) communication is provided by the edge transport nodes of an edge cluster in each location.

The reason to use this method is to avoid managing many tunnels and BFD sessions between all the hosts across locations.

The L2 stretched communication steps are described here:

1. The source ESXi host sends frames to the edge transport node (TEP-to-TEP communication).
2. The source edge transport node forwards a frame to the destination edge transport node (RTEP-to-RTEP communication).
3. The destination edge transport node forwards a frame to the destination ESXi host (TEP-to-TEP communication).

Stretched Layer-2 Network VNI Mapping

When the GM creates a stretched segment, the GM requests each LM to create a local segment. This means that the VNI selection is local to the LM and can be a different VNI. The GM selects the VNI for RTEP. The LM does the TEP-VNI:RTEP-VNI mapping and this mapping is available only in edge transport nodes.

Stretched Layer-2 Packetwalk

The example in Figure 9-27 shows specific edge transport nodes in the Active/Standby mode with segment (Layer-2) stretching.

Virtual Machine 1 on Host Transport Node 1 in Site 1 needs to send a packet to Virtual Machine 2 on Host Transport Node 2 in Site 2.

The steps are described here:

1. Virtual Machine 1 forwards the remote MAC packet to Edge Transport Node 1 (Site 1) over the TEP interface (VNI 4001).
2. Edge Transport Node 1 (Site 1) forwards the packet to Edge Transport Node 1 (Site 2) over RTEP (VNI 6001).
3. Edge Transport Node 1 (Site 2) forwards the packets to Host Transport Node 2 over TEP (VNI 5001).

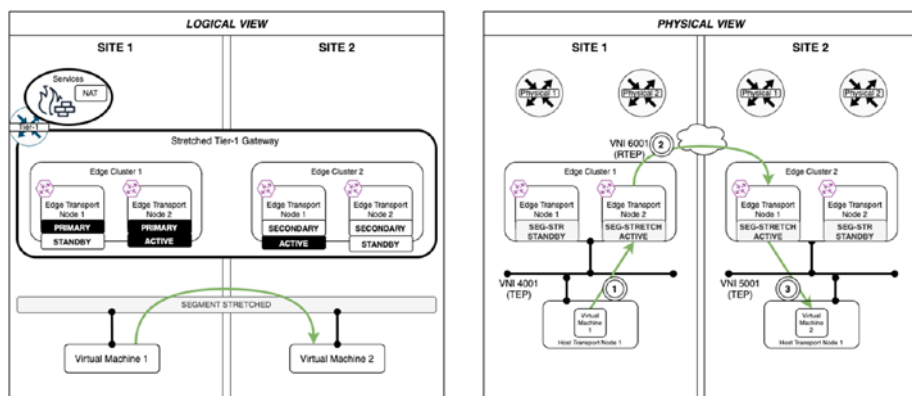


Figure 9-27. Packetwalk diagram

Local Egress Example

Figure 9-28 shows an example of the Layer-3 packet walk, where Site 1 and Site 2 are both configured as primary.

The figure also shows the logical and physical representation of the Tier-0 gateway deployed in Active/Active mode.

The Tier-1 gateway is deployed without services (DR-only). This means that the SR components are not created.

Because both the locations are configured as primary, each location can send northbound egress traffic toward the physical network.

Within each location, the edge transport nodes are in Active/Active configuration, so in Site 1, the Edge Transport Nodes 1 and 2 are active and in Site 2, the Edge Transport Nodes 1 and 2 are active.

Virtual Machine 1 resides on Host Transport Node 1 (Site 1) and Virtual Machine 2 resides on Host Transport Node 2 (Site 2).

Each host transport node sends the traffic to its local edge transport node.

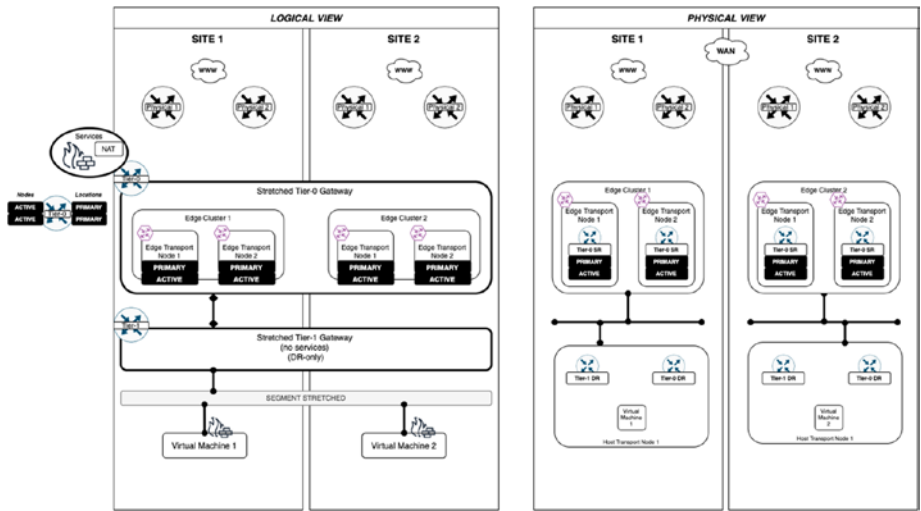


Figure 9-28. Local egress diagram

Virtual Machine 1 on Host Transport Node 1 in Site 1 is trying to access the Internet. The Tier-0 gateway is configured in Active/Active mode, and the Tier-1 gateway is deployed without services (DR-only).

The following steps describe how Virtual Machine 1 will reach the Internet:

1. The traffic from Virtual Machine 1 is sent to the Tier1-DR gateway component on Host Transport Node 1.
2. The Tier1-DR gateway component routes the traffic to the Tier-0-DR gateway component on the same Host Transport Node 1.
3. The Tier-0-DR gateway component sends the traffic to the Tier-0-SR gateway components on Edge Transport Node 1 or Edge Transport Node 2, or both (Site 1).
4. The Tier-0-SR gateway components on Edge Transport Node 1 and Edge Transport Node 2 route the traffic to upstream routers (the Internet).

When there are multiple TCP network streams (egress), the load will be balanced through the Tier-0 SR gateway components on Edge Transport Node 1 and Edge Transport Node 2 (Site 1), but only one single TCP network stream can ingress to any of the single Tier-0 SR gateway components. See Figure 9-29.

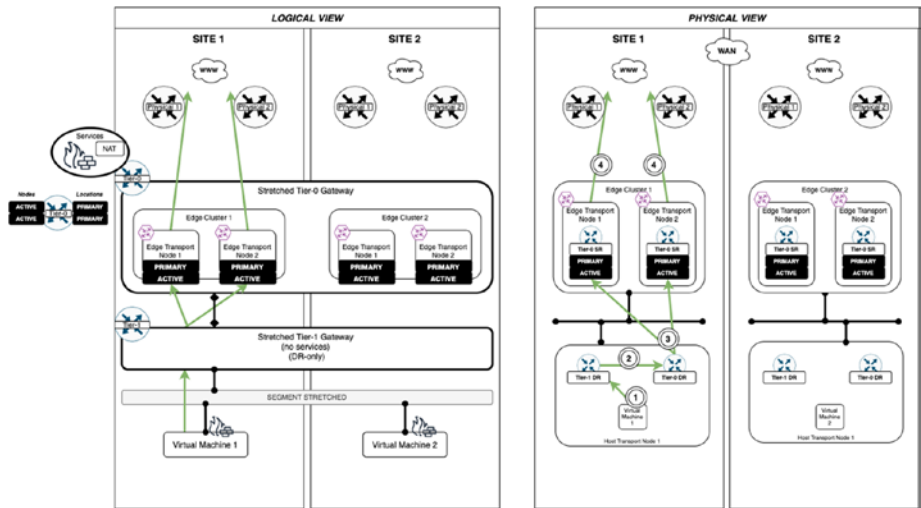


Figure 9-29. Local egress packetwalk

Primary Location Example

Figure 9-30 shows an example where Site 1 is configured as primary and Site 2 is configured as secondary.

The figure also shows the logical and physical representation of the Tier-0 gateway deployed in Active/Active mode.

The Tier-1 gateway is deployed without services (DR-only) and a SR component is not created. In this configuration, only Site 1 can send northbound egress traffic. Within each location, the edge transport nodes are in Active/Active configuration. In Site 1, Edge Transport Nodes 1 and 2 are active and in Site 2, Edge Transport Nodes 1 and 2 are also active.

Virtual Machine 1 resides on Host Transport Node 1 located in Site 1, and Virtual Machine 2 resides on Host Transport Node 2 located in Site 2.

Each ESXi node sends the traffic to its local edge transport node.

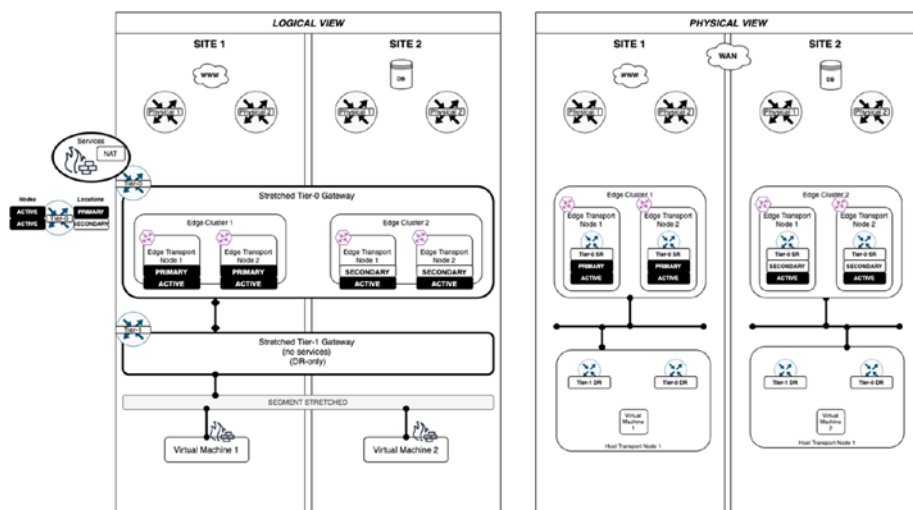


Figure 9-30. Primary location diagram

Now let's assume that Virtual Machine 1 in Site 1 needs to communicate with the Internet.

The steps for Virtual Machine 1 to reach the Internet are listed here:

1. The traffic from Virtual Machine 1 is sent to the Tier-1-DR gateway component on Host Transport Node 1.
2. The Tier-1-DR gateway component routes the traffic to the Tier-0-DR gateway component on the same Host Transport Node 1.
3. Tier-1-DR gateway component sends the traffic to the Tier-0-SR gateway components on Edge Transport Node 1 or Edge Transport Node 2, or both (Site 1).
4. The Tier-0-SR gateway components on Edge Transport Node 1 and Edge Transport Node 2 (Site 1) route the traffic to upstream routers (the Internet).

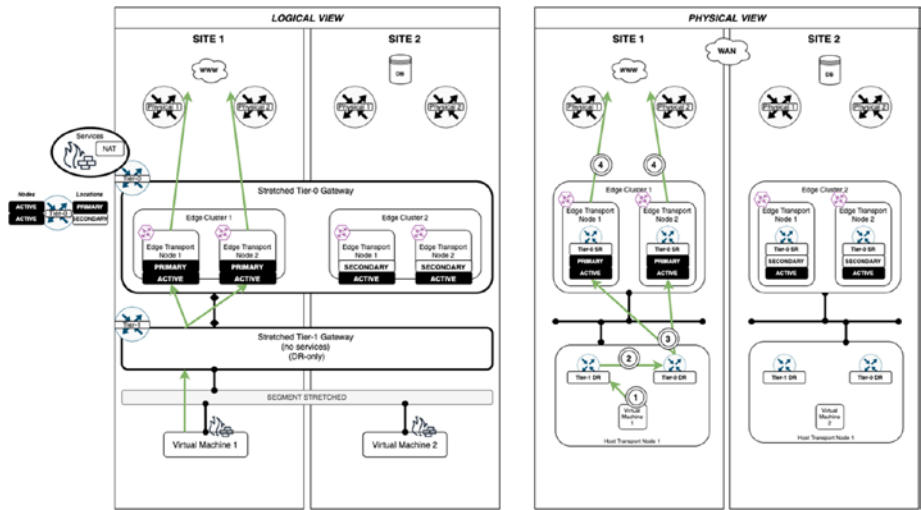


Figure 9-31. Primary location packetwalk (Internet access) (1)

Now let's assume that Virtual Machine 2 in Site 2 needs to communicate with the Internet.

The steps for Virtual Machine 2 to reach the Internet are listed here:

1. The traffic from Virtual Machine 2 is sent to the Tier-1-DR gateway component on Host Transport Node 2.
2. The Tier-1-DR gateway component routes the traffic to the Tier-0-DR gateway component on the same Host Transport Node 2.
3. The Tier-0-DR gateway component sends the traffic to the Tier-0-SR gateway components on Edge Transport Node 1 or Edge Transport Node 2 (Site 2), or both.

4. The Tier-0-SR gateway components of Edge Transport Node 1 and Edge Transport Node 2 forward the traffic to the Tier-0-SR gateway components situated in Edge Transport Nodes 1 and 2 (Site 1).
5. The Tier-0-SR gateway components on Edge Transport Node 1 and Edge Transport Node 2 (Site 1) route the traffic to upstream routers (the Internet). See Figure 9-32.

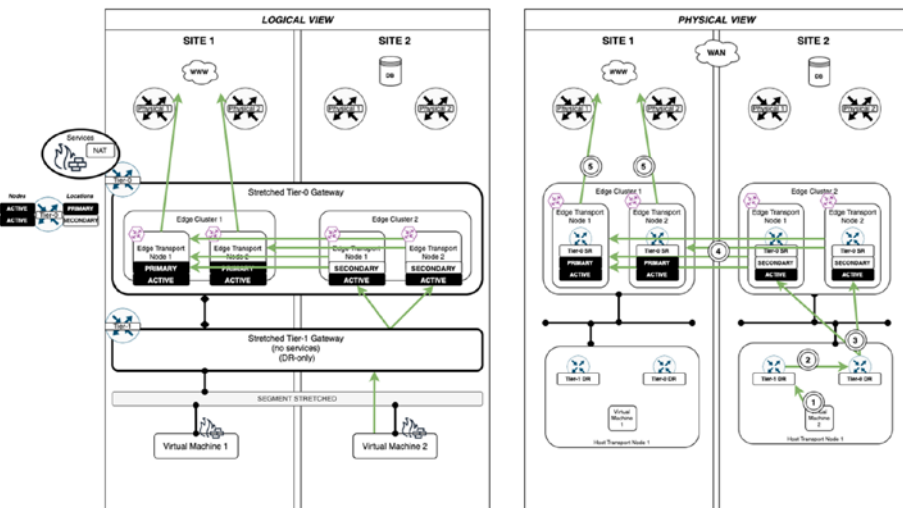


Figure 9-32. Primary location packetwalk (Internet access) (2)

Now let's assume that Virtual Machine 1 at Site 1 needs to communicate with the database server in Site 2.

The steps for Virtual Machine 1 (Site 1) to reach the database server (Site 2) are listed here:

1. The traffic from Virtual Machine 1 is sent to the Tier-1-DR gateway component on Host Transport Node 1.
2. The Tier-1-DR gateway component routes the traffic to the Tier-0-DR gateway component on the same Host Transport Node 1.
3. The Tier-0-DR gateway component sends the traffic to the Tier-0-SR gateway components on Edge Transport Node 1 or Edge Transport Node 2, or both (Site 1).
4. The Tier-0-SR gateway components of Edge Transport Nodes 1 and 2 (Site 1) forward the traffic to the Tier-0-SR gateway components situated in Edge Transport Nodes 1 and 2 on the other side (Site 2).
5. The Tier-0-SR gateway components on Edge Transport Nodes 1 and 2 (Site 2) route the traffic to upstream routers to reach the destination (the database server). See Figure 9-33.

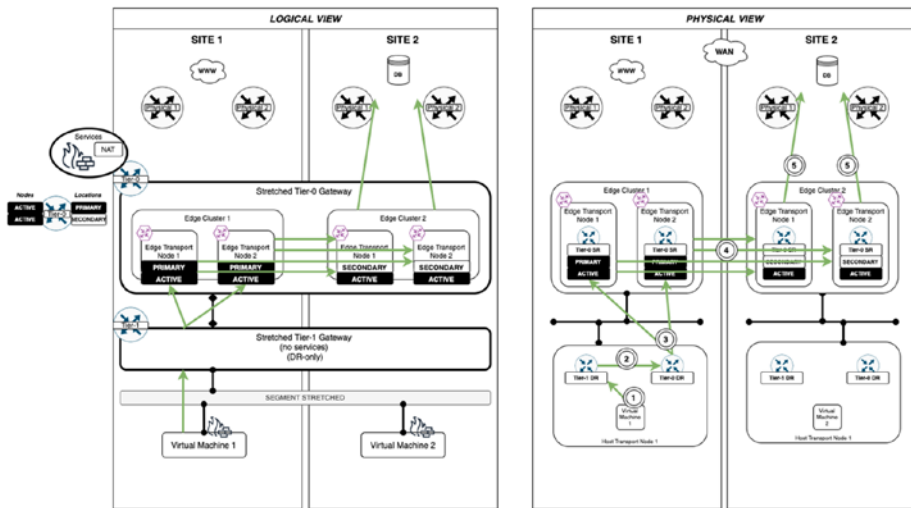


Figure 9-33. Primary location packetwalk (Database access) (1)

In this last scenario, let’s assume that Virtual Machine 2 in Site 2 needs to communicate with the database server (Site 2).

The steps for Virtual Machine 2 (Site 2) to reach the database server (Site 2) are listed here:

1. The traffic from Virtual Machine 2 is sent to the Tier-1-DR gateway component on Host Transport Node 2.
2. The Tier-1-DR gateway component routes the traffic to the Tier-0-DR gateway component on the same Host Transport Node 2.
3. The Tier-0-DR gateway component sends the traffic to the Tier-0-SR gateway components on Edge Transport Node 1 or Edge Transport Node 2 (Site 2), or both.

4. The Tier-0-SR gateway components on Edge Transport Nodes 1 and 2 (Site 2) route the traffic to the upstream routers to reach the destination (the database server). See Figure 9-34.

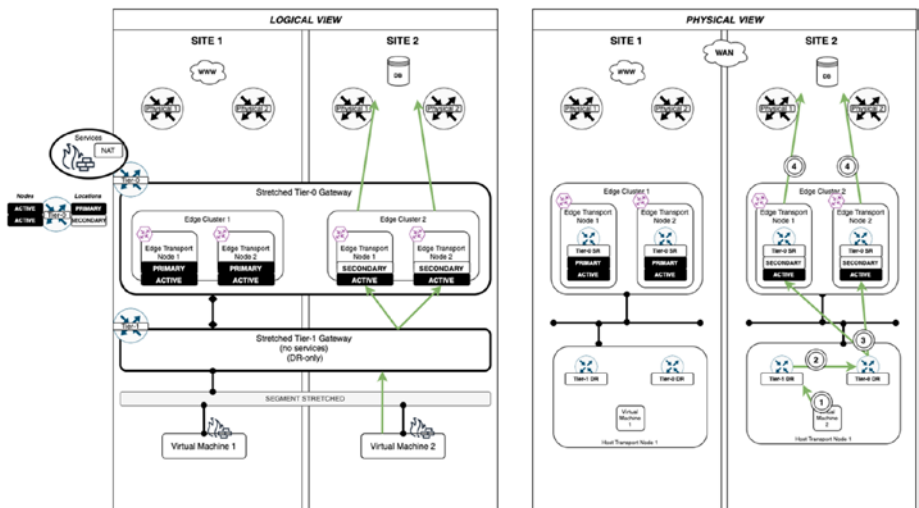


Figure 9-34. Primary location packetwalk (database access) (2)

NSX-T Federation Security

This section explains the NSX-T Federation security use cases and describes the Federation security components with the security configuration workflows related to NSX-T Federation.

NSX-T Federation Security Use Cases

NSX-T Federation provides central security policy for data centers managed on-premises. This works for use cases such as data center extensions with a centralized pane, with the goal to provide consistency. NSX-T Federation extends to multiple sites/data centers.

NSX-T Federation Security Components

NSX-T Federation security uses the components outlined in Table 9-3.

Table 9-3. *NSX-T Federation Security Components*

Component	Description
Rules	The firewall rules are based on objects and tags instead of IP addresses.
Policy	The firewall policy includes one or more individual firewall rules.
Global groups	The groups created by the GM.
Local groups	The groups created by the LM.
Tags	The objects that can be tagged and searched by a specific location or can be defined globally and added to a global group.
Regions	The region is a collection of locations. Each location that is added to the GM becomes a region by default. You can create custom regions if you want.

GM Firewall Policy

When you configure a security policy, you can use the settings listed in Table 9-4.

Table 9-4. *GM Security Policy “Applied To” Specifics*

Applied To in the policy is set to DFW	All logical switch ports (virtual machines and containers) of the span receive the rules in that section.
Applied To in the policy is set to Group	All group members (virtual machines and containers) receive the rules in that section.

You can create a global security policy by navigating using the GM to Security ► East West Security ► Distributed Firewall. See Figure 9-35.

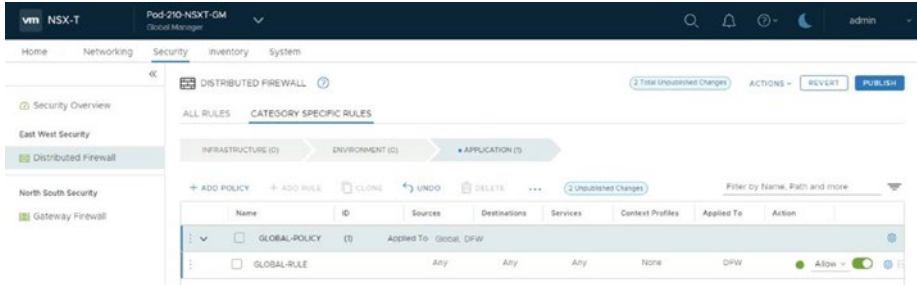


Figure 9-35. Global Security Policy

GM Firewall Rules

When you configure a firewall rule, use the settings listed in Table 9-5.

Table 9-5. GM Firewall Rules “Applied To” Specifics

Applied To in the policy is set to DFW	All logical switch ports (virtual machines and containers) of the span receive the rules within that section.
Applied To in the policy is set to Group	All group members (virtual machines and containers) receive the rules within that section.

You can create a global firewall rule by navigating using the GM to Security ► East West Security ► Distributed Firewall. See Figure 9-36.

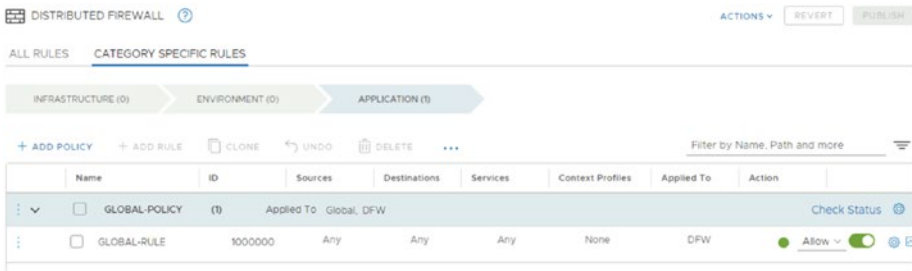


Figure 9-36. Global Firewall Rule

GM and LM Section Overlap

The GM and LM rules are displayed in a mixed way when you look at the LM Distributed Firewall section.

For GM and LM sections created in the same category, the GM sections are always at the top, as you can see in Figure 9-37.

You can see this by navigating using the LM to Security > East West Security > Distributed Firewall. See Figure 9-37.

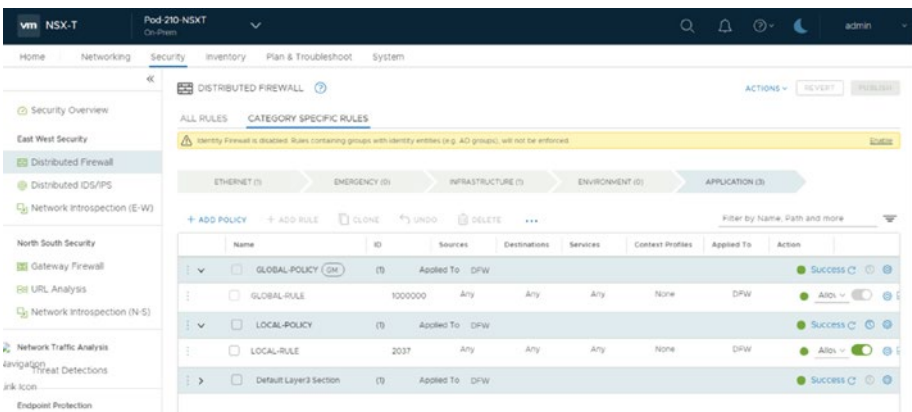


Figure 9-37. GM and LM firewall rules overlap

The gateway firewall rules can also be created in the same category.

The GM and LM sections are individually created under the T0/T1 gateway firewall and under the distributed firewall. The GM rules are always above the LM.

Global Groups

The GM creates global groups. The global group types are either members owned by NSX-T or discovered members.

Members that are owned by NSX-T are GM-created groups or objects. Discovered members are LM-created objects that GM discovers on demand and tags later.

From the LM, the GM groups are visible. The LM groups are not visible from the GM, but can be requested.

You can use the GM to create, update, and delete global groups.

Both the GM and LM can read global groups, but global groups cannot be used as nested members in LM-based policies.

Global Groups GUI

You can create a global group by navigating using the GM to Inventory ► Groups. See Figure 9-38.

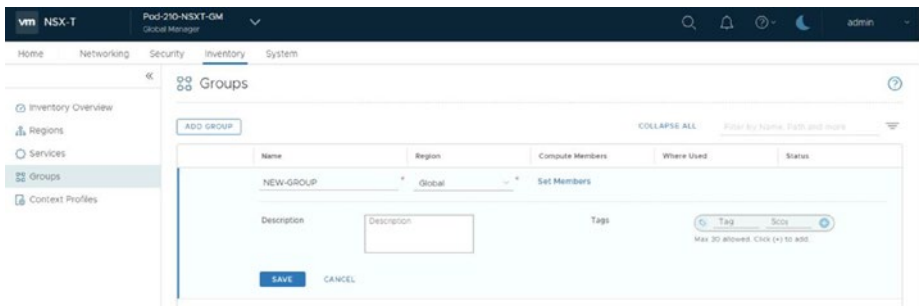


Figure 9-38. Global groups

Regions

A region creates a focused group for security and networking policies.

Some regions are created automatically by default after the onboarding process in the GM. This is typically the site or location you have just onboarded, but if you want to have a subset of multiple regions, you need to create this region manually.

You can create a region by navigating using the GM to Inventory ► Regions.

In Figure 9-39, you see the regions that were configured by default when I onboarded the Federation sites to the GM.

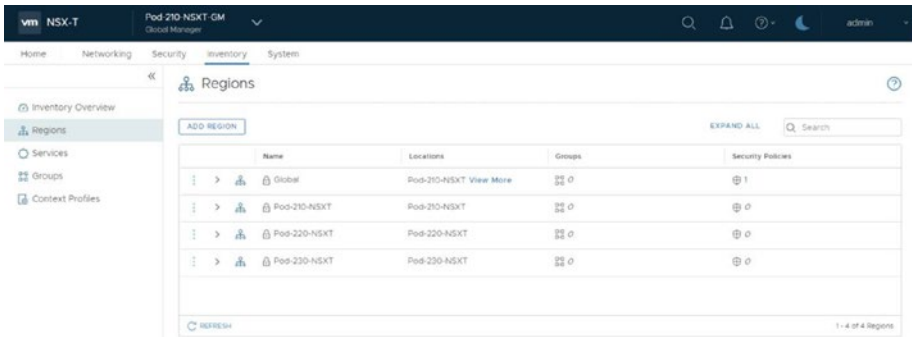


Figure 9-39. *Regions*

In Figure 9-40, I manually created a region that consists of two sites.

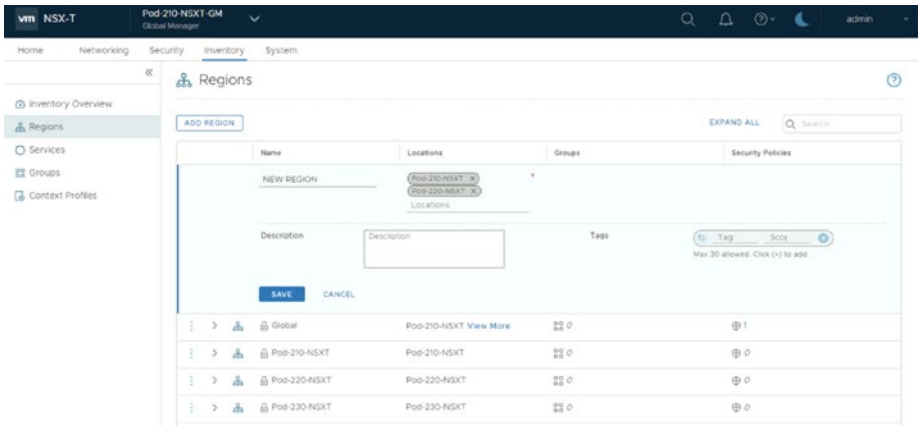


Figure 9-40. Manually created region with two sites (1)

In Figure 9-41, you see the result of the region that I created manually.

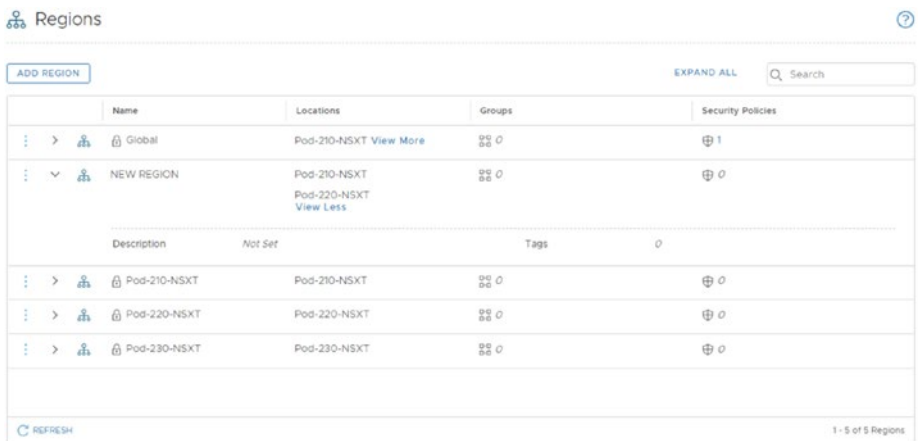


Figure 9-41. Manually created region with two sites (2)

NSX-T Federation Security Tags

With the tag-based criteria, you can select local and global objects.

Because the grouping behavior does not differentiate between local and global members, tag-based criteria can include local and global objects.

You can set the criteria inside the group. You can select the members and then select a tag to make an object a member of that group. See Figure 9-42.




Figure 9-42. Group membership criteria

Federation Security Configuration Steps

The overall architectural steps (under the hood) that NSX-T Federation takes are described here (Figure 9-43):

1. You send the configuration to the GM.
2. The GM then sends the configuration to the LM for each location.
3. The LM then forwards the configuration to the management plane/central control plane.
4. The central control planes from each location then synchronize the configuration across all locations.
5. Each central control plane pushes the consolidated configuration to the data plane (local for each site).

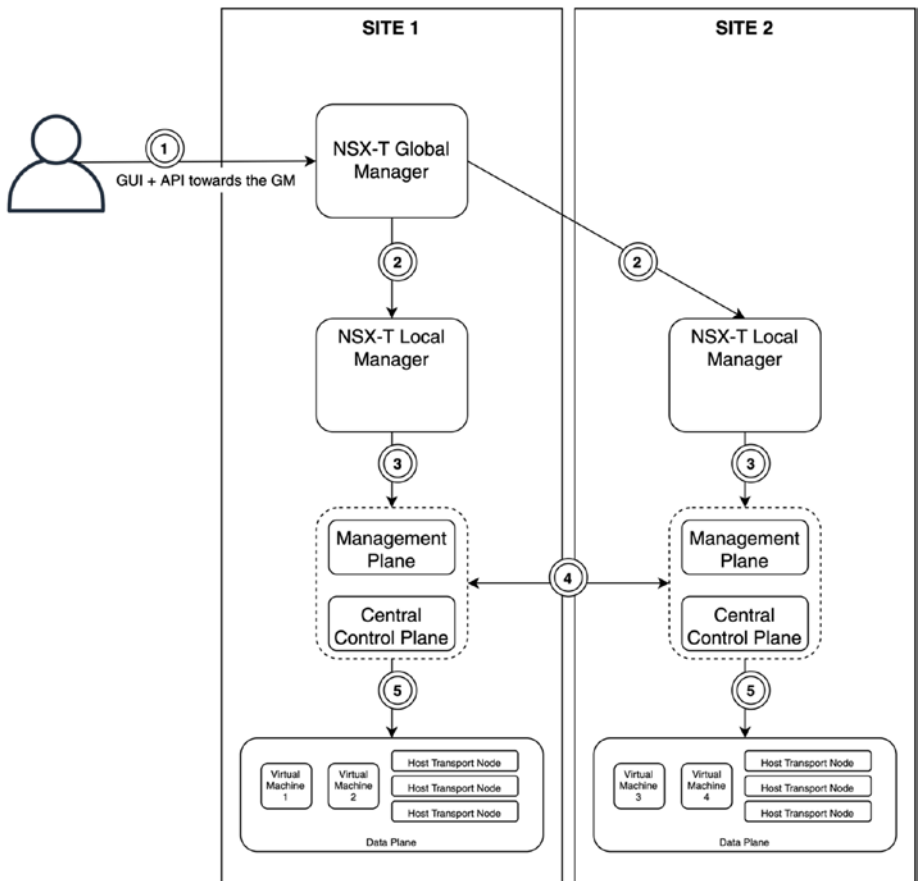


Figure 9-43. NSX-T configuration steps

GM Group Types and Span of Group Types

GM security groups fall into three types: *Global*, where all sites are stretched, *Region*, where a subset of the sites are stretched, and *Local*, where the group is not stretched but is local to a specific site.

All these groups are created from the GM. See Figure 9-44.

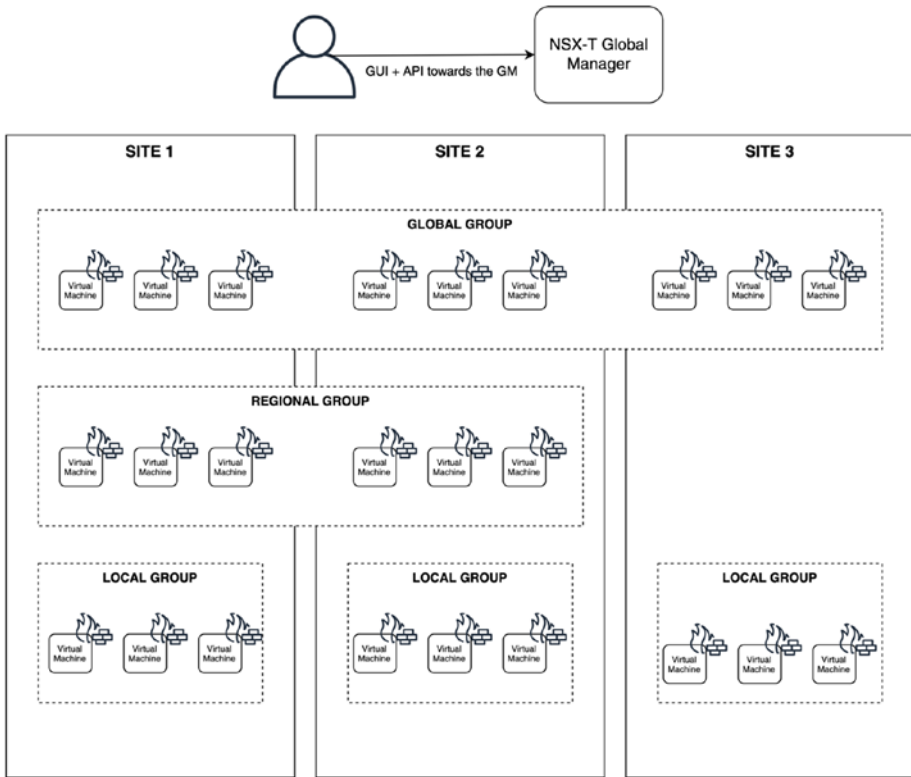


Figure 9-44. GM security groups (global, regional, and local)

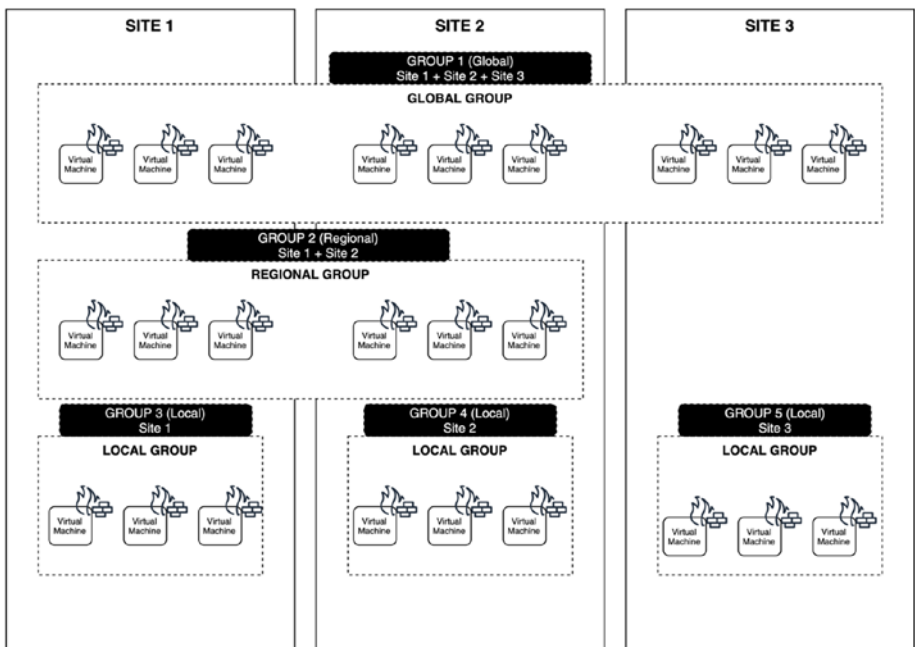
The GM groups span is Global, Regional, or Local. The membership is Static or Dynamic.

The GM pushes groups to each LM in the span. For dynamic groups, each LM resolves its local members and updates (with the local objects it has to offer) other LMs that belong to the span. See Table 9-6.

Figure 9-45 shows a span of the groups.

Table 9-6. *Span of Groups*

Group 1	Global and stretched to all sites, that is, Sites 1, 2, and 3.
Group 2	Region and only stretched to Site 1 and 2.
Group 3	Only with Site 1 and local to Site 1.
Group 4	Only with Site 2 and local to Site 2.
Group 5	Only with Site 3 and local to Site 3.

**Figure 9-45.** *Span of groups*

GM Group Rules

As you have seen, groups can be global or local. If a policy scope is local (for example, Site 1), the rules inside the policy should be limited to the same scope (Site 1). If a policy is global, the rules should be global (Sites 1, 2, and 3).

Figure 9-46 shows an example where I configured firewall rules between a global group and a regional group, but there are also firewall rules between the two local groups.

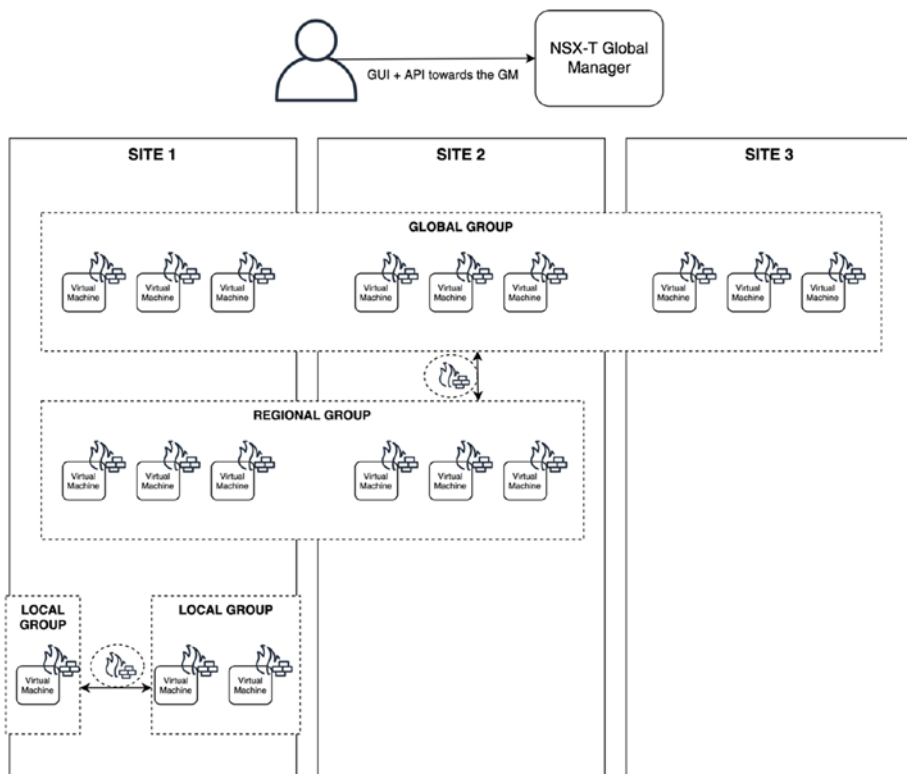


Figure 9-46. GM group firewall rules

GM Group Span of Dynamic Members

The span of the group should always be more than the span of the objects that you associated with it.

In Figure 9-47, Group 2 contains a segment connected to the stretched Tier-1 gateway. This scope (stretched segment) has a larger scope than Group 2. This means Virtual Machine 3 will not contain the required policies (because it is not part of the group).

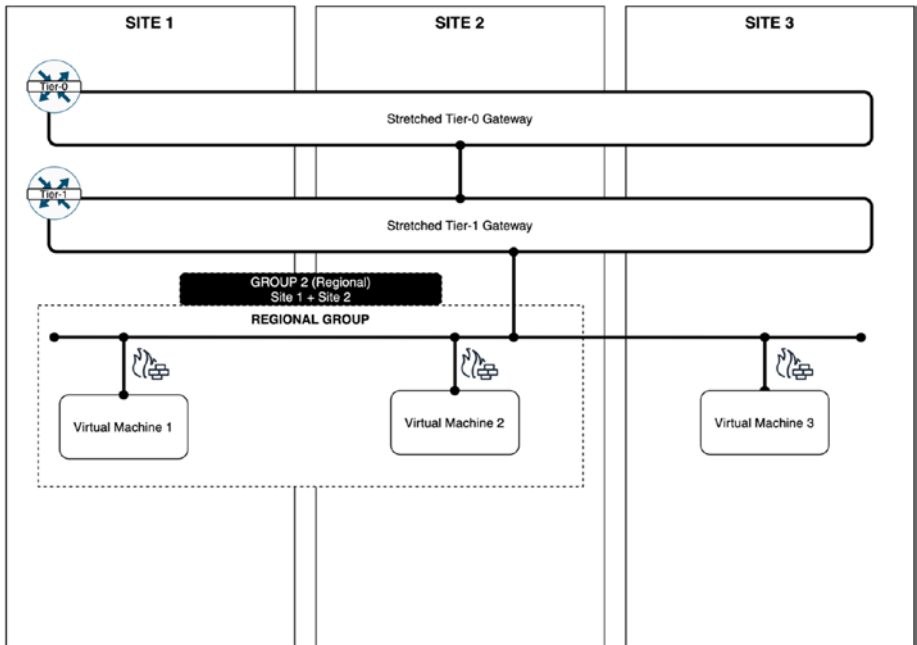


Figure 9-47. Group span with dynamic group members

GM Group Based on a Virtual Machine Tag

Dynamic groups can contain virtual machine-based tags.

Figure 9-48 shows groups that are associated with virtual machine-based tags for dynamic membership. This will result in Virtual Machines 1, 2, and 3 all becoming part of the group.

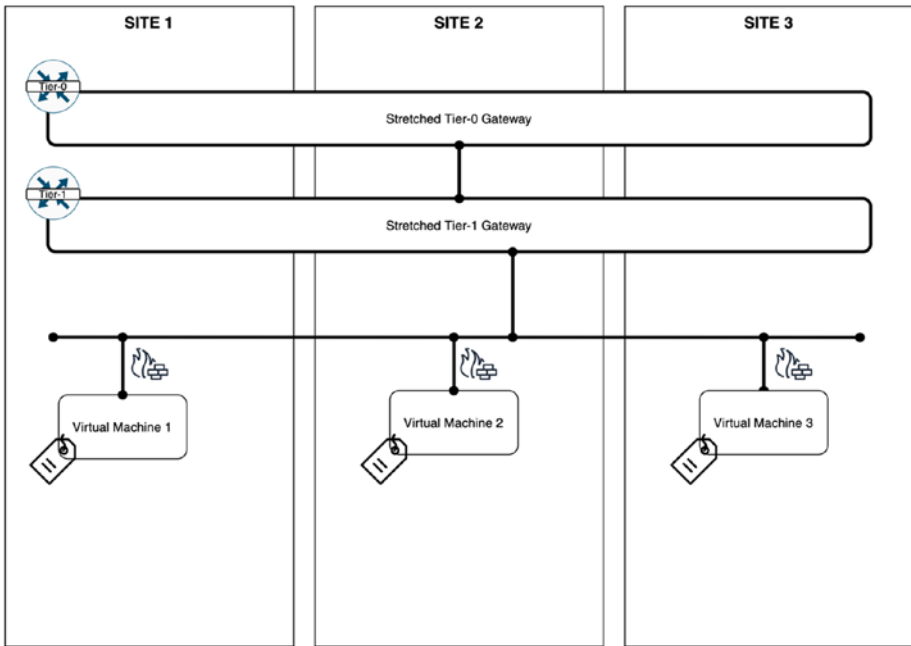


Figure 9-48. Group span with dynamic group members based on tags

Tags are applied to virtual machines even after the virtual machine is moved from one site to another using vMotion. This means that the group rules (and therefore the security policy) are intact. See Figure 9-49.

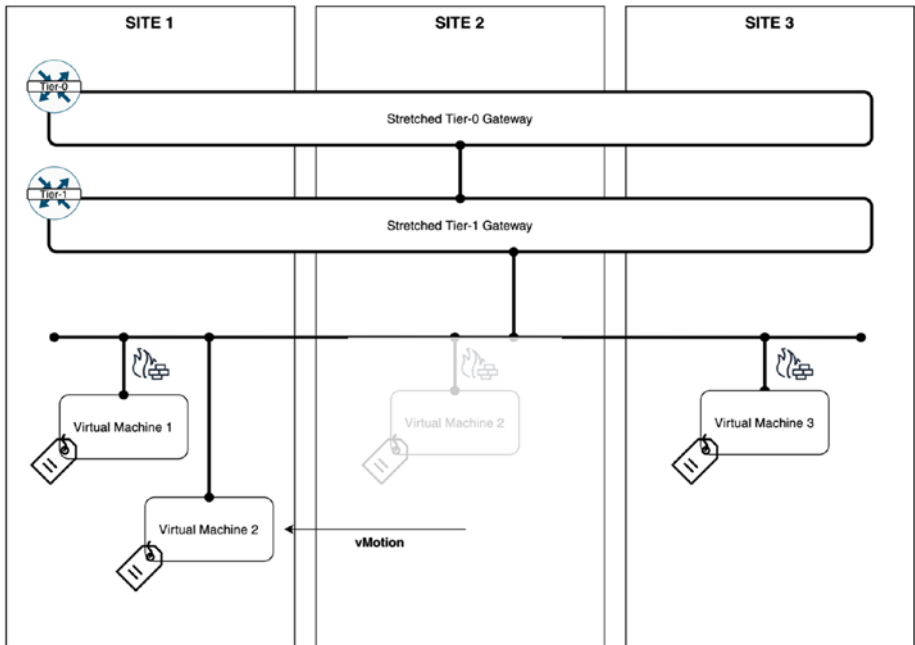


Figure 9-49. Group span with dynamic group members based on tags with security policy retention

Summary

This chapter covered the NSX-T Federation concept, including its networking and security capabilities. NSX-T Federation provides network and security services stretched across multiple sites.

The next chapter discusses how NSX-T is used inside the commonly used VMware public cloud offerings (hyperscalers).

CHAPTER 10

Public Cloud Integration

This chapter explains how NSX-T is used inside the commonly used VMware public cloud offerings (hyperscalers) and the commonly used native public cloud providers.

NSX-T on VMware Hyperscalers

This section explains what a VMware hyperscaler is and how NSX-T is used and implemented from a high level on the commonly used VMware public cloud offering.

What Is a VMware Hyperscaler?

A hyperscaler from a VMware perspective is a public cloud environment in which VMware VCF is deployed using the infrastructure of a native public cloud provider.

The benefit of this offering is that a complete VMware VCF environment that manages your on-premises locations is now available in the public cloud.

This makes it easier to extend your applications to the public cloud using the same VMware virtual machine workload constructs. See Figure 10-1.





Figure 10-1. Hyperscaler

VMware has partnered with some cloud providers to offer the “managed VMware SDDC” service.

The current largest public cloud providers that VMware has partnered with are AWS, Microsoft, and Google. Each cloud provider uses its own branded name (Figure 10-2). Table 10-1 shows the naming/branding of the largest public cloud providers for the hyperscalers.

Table 10-1. *Public Cloud Providers for the Hyperscalers*

Public Cloud Provider	Hyperscaler Name
Google	Google Cloud VMware Engine (GCVE)
AWS	VMware Cloud on AWS (VMC on AWS)
Microsoft	Azure VMware Solution (AVS)

**Figure 10-2.** *Public cloud providers for the hyperscalers*

VMware’s current hyperscaler partners can be found on this website (Figure 10-3):

<https://www.vmware.com/asean/partners/work-with-partners/hyperscalers.html>



Figure 10-3. Hyperscaler partners of VMware

VMware Hyperscaler Use Cases

Use cases for a VMware hyperscaler are described in Table 10-2.

Table 10-2. VMware Hyperscale Use Cases

Use Case	Description
Data center expansion/ capacity burst/ maintenance	A VMware hyperscaler can be used if you need additional resources and you need them fast. Additional resources can be desired when you suddenly need to expand, or if you need a place to host your virtual machines to keep your applications up and perform maintenance on your existing on-premises environment.

(continued)

Table 10-2. *(continued)*

Use Case	Description
Application proximity	You can use a VMware hyperscaler when you want to offer your applications inside a specific country. When you keep the application access proximity close to a specific country, you may improve performance and maybe even save money.
Data center migration	A VMware hyperscaler can also be used when you want to decommission one or multiple on-premises data centers, and you need to migrate your application virtual machine workloads elsewhere.
Disaster recovery	You can also use a VMware hyperscaler as your disaster recovery (DR) site if you have any failures on-premises.

NSX-T Deployment Footprint

The NSX-T deployments inside the different VMware hyperscaler offerings of GCVE, VMC on AWS, and AVS and their differences are described in Table 10-3.

Table 10-3. *NSX-T Deployment Footprint*

	GCVE	VMC on AWS	AVS
NSX-T Managers	Three NSX-T Manager nodes deployed		
NSX-T Edge Transport Nodes	Two NSX-T Edge transport nodes (VMs) deployed		
NSX-T Host Transport Nodes	This is either three or the maximum number of hosts allowed in the vSphere cluster		
Connectivity methods between on-premises and off-premises	Interconnect	Direct connect	ExpressRoute
NSX-T Management	Direct access to the NSX-T Manager	VMC console	Direct access to the NSX-T Manager
NSX-T release	This can be different for each hyperscaler and is determined by the cloud provider		
Native cloud routers/gateways	Cloud router	TGW/VGW/DGW	VNet gateway
NSX-T gateways	By default, one Tier-0 gateway and one Tier-1 gateway		

Note You do *not* have control over the NSX-T deployment footprint in VMware hyperscaler environments; the NSX-T infrastructure lifecycle management is completely handled by VMware.

NSX Cloud

This section explains what NSX Cloud is and discusses its main use cases.

NSX Cloud Use Cases

NSX Cloud offers consistent networking and security for applications that are running on your on-premises data center and in the native public cloud.

When I say “native” public cloud, I am not referring to the VMware hyperscalers, but to the native AWS or Azure cloud environments.

Both AWS and Azure have their own way of configuring network and security features. The objects they use to offer network services are also completely different.

NSX Cloud places an additional layer on top so that networking and security will have the same look and feel that you are used to in NSX-T from an operational and management perspective.

In order to do this, some additional components are deployed on-premises and inside the native public cloud constructs.

Note At the time of this writing, only AWS and Azure are supported by NSX Cloud. From Azure, you can deploy the entire stack within a subscription without the need to have anything on-premises. (NSX-T Manager cluster and CSM are deployed in the resource group.)

NSX Cloud Components

To deploy NSX Cloud, you need to run NSX-T on-premises and deploy a cloud services manager (CSM) and a public cloud gateway (PCG). See Figure 10-4.

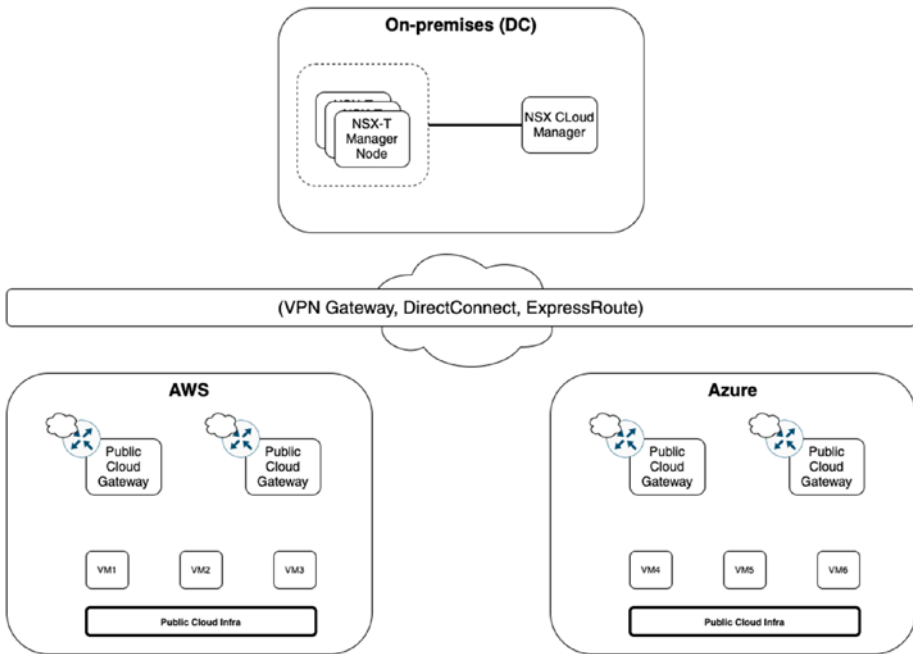


Figure 10-4. NSX Cloud components

NSX-T Manager (Cluster)

The NSX-T Manager (cluster) is something you typically would already have running on-premises. If not, this is a prerequisite before you can deploy NSX Cloud.

NSX Cloud Service Manager

The purpose of the NSX Cloud Service Manager (CSM) is to provide a unified view between the running NSX-T state and the public cloud inventory. You will be able to onboard new VPCs (AWS) or VNets (Azure) by automating the public cloud gateway deployment. You can also configure a default quarantine policy and perform lifecycle management tasks on the public cloud gateway and NSX agents.

There is a one-to-one mapping between the NSX Cloud Service Manager and the NSX-T Manager.

NSX Public Cloud Gateway

The NSX public cloud gateway (PCG) provides north-south connectivity between the public cloud and the on-premises management components of the NSX-T data center.

The purpose of the PCG is to act as the local NSX control plane inside the public cloud. It performs an inventory discovery within VPC (AWS) or VNET (Azure) and it is responsible for configuring and adopting cloud tags.

With the PCG, you can also enforce a default quarantine policy and the PCG will act as an NSX edge gateway for north-south traffic when in the path (using NSX Enforced Mode only).

The PCG appliance is deployed in HA (Active-Standby) and is available as a public AMI in AWS or copied over as a VHD in Azure.

The preferred location to deploy the PCG is in the transit VPC (AWS) or VNET (Azure), but it can also be deployed in a regular VPC or VNET.

Multi-Cloud Deployment

There are two deployment architectures in which you can use NSX Cloud—direct connectivity or a transit network with peering.

Deployment Architectures: Direct Connectivity

Figure 10-5 shows a deployment architecture that can apply AWS and Azure. The PCGs are deployed in the VPC/VNET and the workloads leverage the services offered by these PCGs.



Figure 10-5. *Deployment architectures: direct*

Deployment Architectures: Transit Network with Peering

Figure 10-6 shows a deployment architecture that can apply AWS and Azure. This time, the PCGs are deployed in the transit VPC/VNET and the workloads can be in another VPC/VNET dedicated to compute. That way, they leverage the services offered by these PCGs from a central point.

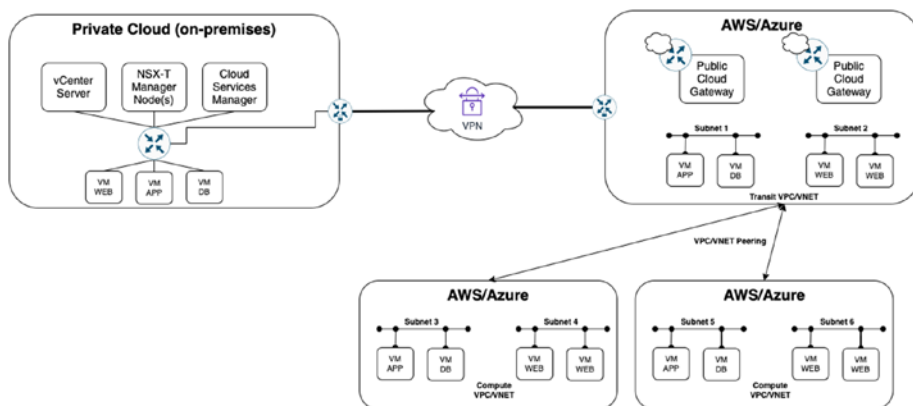


Figure 10-6. *Deployment architectures: transit network*

Deployment Modes

The virtual machines inside the public cloud can use the network and security services offered by the public cloud gateway through two deployment models—cloud enforcement mode and NSX Tools mode.

Cloud Enforcement Mode

Cloud enforcement mode provides you with a common policy framework. NSX-T translates NSX policies to native cloud specific security policies. Translated security policies are subject to any cloud provider limits (scale, feature, etc.).

In this mode, no NSX Tools are deployed on the virtual machines. The granularity of control is at the VPC/VNET level and everything inside the managed VPC/VNET is managed unless it's whitelisted.

The default quarantine policy does not apply in this mode.

Essentially, a replica of the security policy rules (configured in the NSX-T Manager) is not translated and programmed into AWS or Azure, based on their own security policy group and rule constructs.

NSX Tools Mode

NSX Tools mode provides a consistent policy framework where NSX policies are enforced within NSX Tools. The cloud provider limits do not apply here.

NSX Tools must be deployed on the virtual machines. The granularity of control is at each individual virtual machine level. Only tagged virtual machines are “managed”.

The default quarantine policy provides an additional layer of defense and whitelisting should be done through the CSM.

Essentially, all traffic coming from the public cloud native virtual machine is redirected to the PCG. The PCG uses the same security construct as was enforced by the NSX-T Manager.

Deployment Steps

This section outlines the high-level deployment steps required to deploy NSX Cloud in an AWS and Azure public cloud. These steps will give you an overview and a good idea as to what is involved.

On AWS

1. Get a native AWS account.
2. Deploy the NSX-T Manager node(s).
3. Assign static IP addresses to the NSX Management nodes and the NSX Cloud Service Manager.
4. Configure VIP for the NSX management cluster.
5. Configure Forward and Reverse records for the NSX Manager VIP, nodes, and the NSX Cloud Service Manager in your internal DNS domain.
6. Deploy a CSM.
7. Integrate the NSX-T with the CSM.
8. Establish connectivity between the on-premises and AWS VPC.
9. Deploy the NSX public cloud gateway in AWS VPC.
10. Generate NSX Cloud permission in AWS.
11. Add the AWS account details in CSM.
12. Deploy a test virtual machine using the NSX Tools.
13. Deploy a test virtual machine using native cloud enforcement.
14. Create a security policy on the NSX-T Manager and verify if the policy is applied to the virtual machine with NSX Tools and the native cloud enforcement.
15. Test the routing with the virtual machine using the NSX Tools.

On Azure

1. Get a native Azure account.
2. Deploy the NSX-T Manager node(s).
3. Assign static IP addresses to the NSX Management nodes and the NSX Cloud Service Manager.
4. Configure VIP for the NSX management cluster.
5. Configure Forward and Reverse records for the NSX Manager VIP, nodes, and the NSX Cloud Service Manager in your internal DNS domain.
6. Deploy a CSM (when deployed for AWS, use the same CSM).
7. Integrate NSX-T with the CSM (when deployed for AWS, use the same CSM).
8. Establish connectivity between the on-premises and VNET domains.
9. Deploy the NSX public cloud gateway in Azure VNET.
10. Generate NSX Cloud permissions in Azure.
11. Generate an Azure Service Principal.
12. Add Azure subscription details to the CSM.
13. Deploy a test virtual machine using the NSX Tools.
14. Deploy a test virtual machine using native cloud enforcement.

15. Create a security policy on the NSX-T Manager and verify if the policy is applied to the virtual machine with NSX Tools and the native cloud enforcement.
16. Test the routing with the virtual machine using the NSX Tools.

Summary

This chapter explained what a VMware hyperscaler is and how NSX-T is positioned inside the different VMware hyperscalers. It also explained what NSX Cloud is and discussed its use cases and capabilities in combination with the AWS and Azure native public cloud providers.

The next chapter covers cloud management platform integration and cloud automation.

CHAPTER 11

Cloud Management Platform Integration and Automation

NSX-T's REST API can be consumed by any application. That means it's possible to perform Create, Read, Update, and Delete (CRUD) operations. There are several VMware products and other products (such as API wrappers) that use the NSX-T API to automate certain tasks. This chapter briefly explains some of these products and capabilities. As cloud management platform integration and cloud automation can be very complex topics and complete books can be written on them, I only touch the surface here.

VMware Cloud Management Platforms

VMware has its own set of products that can integrate with NSX-T in order to perform CRUD operations. vRealize Automation, vRealize Orchestrator, and Cloud Director are examples of these products. This section briefly explains these products and discusses how they integrate with NSX-T.

vRealize Automation (vRA)

vRA is also known as vRealize Automation; it allows you to design “cloud templates” (in the past these were called *blueprints*). In these templates, you can create or use a subset of NSX-T objects/features.

In Figure 11-1, you see an example of a cloud template where I am deploying two virtual machines, a WEB, and a DB virtual machine, and I am placing these on an existing NSX-T segment (a network).

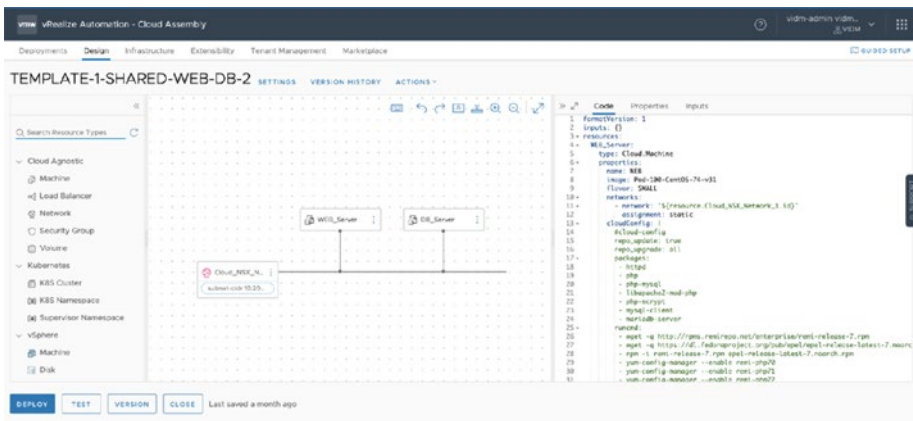


Figure 11-1. vRA cloud template

There are many more out-of-the-box NSX-T features and objects you can use within a blueprint.

A good use case is if you want to automate the deployment of an application that has multiple tiers. This application needs to be attached to specific networks and needs to adhere to the existing security policy. Using vRA and NSX-T, you can develop a solution in a few hours and roll out a complete multitiered application in minutes.

Whenever there is a task that you want to automate inside NSX-T, but it’s not possible out-of-the-box using vRA, you can do it with vRO.

vRealize Orchestrator (vRO)

vRO is also known as vRealize Orchestrator; it's capable of performing complex automation tasks (not only on NSX-T) with the use of complex workflows. Many integrations are possible and the possibilities are almost endless when you start working with vRO. With vRO you can execute REST API calls to the NSX-T Manager in order to perform certain tasks that you can't do using vRA.

VMware Cloud Director

VMware Cloud Director is typically used by service providers that have multiple tenants and need to provide network and security capabilities to these tenants. The tenants usually have limited rights to create a set of virtual machines (vApp) consisting of an application. The underlying infrastructure is still one NSX-T Manager installation. The cloud director, integrated into NSX-T, maintains multi-tenancy by carving out gateways and networks so that they can only be presented to a designated tenant.

Figure 11-2 shows an example of a single provider VDC with a single Tier-0 gateway (with VRF enabled Tier-0 gateways). It also shows multiple organization VDCs, each having a dedicated Tier-1 gateway (edge) with dedicated networks assigned.

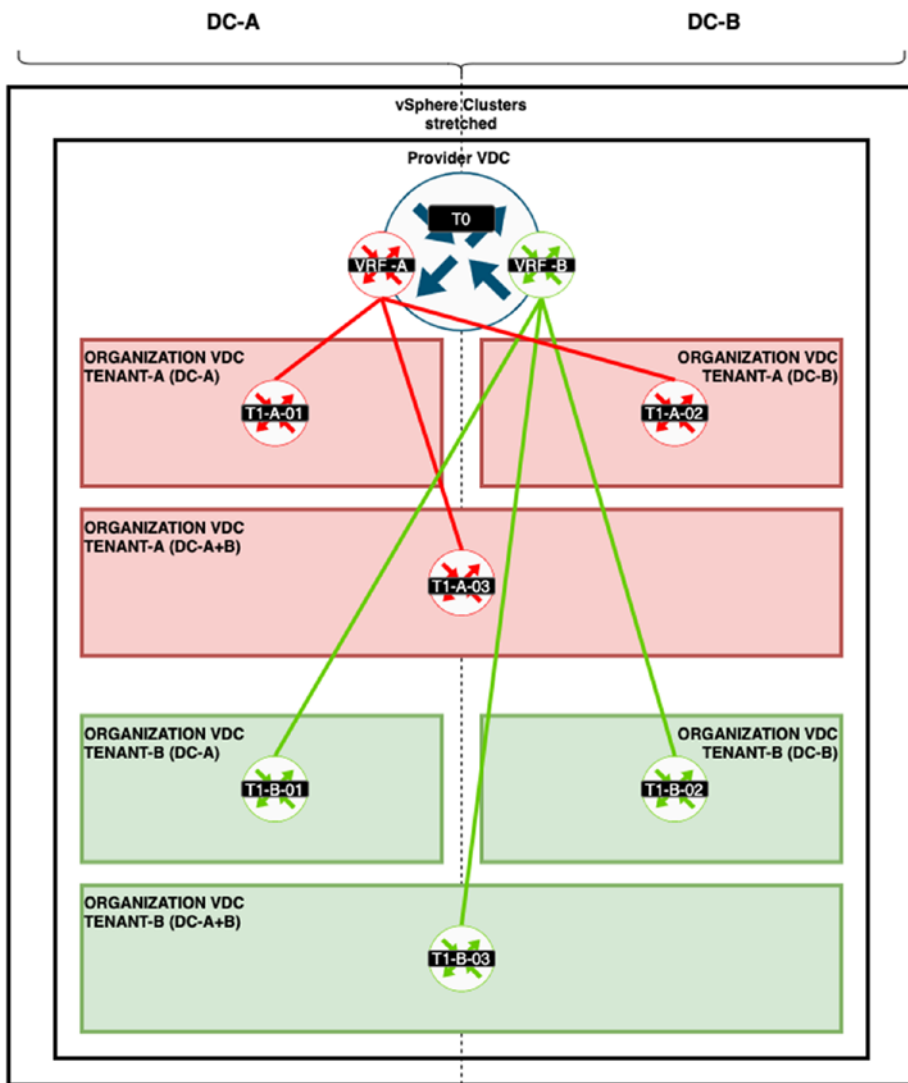


Figure 11-2. Single provider VDC with a single Tier-0 gateway

You can find a detailed deployment of the Cloud Director integrated with NSX-T at https://nsx.ninja/index.php/VMware_Cloud_Director_Basics_with_NSX-T.

Other Cloud Management Platforms

Besides the VMware branded products, there are also some good tools to consider. Ansible, Terraform, and PowerShell/PowerCLI are all very important if you are looking at Infrastructure as Code (IaC) and can deploy and configure complete topologies and sometimes can even deploy the actual underlying infrastructure.

Note Operations that are executed using API directly or as code can have a catastrophic impact on the infrastructure if they're not properly tested and evaluated.

Ansible

Ansible is an API wrapper that uses special NSX-T “modules” in combination with YAML code to perform certain automation tasks. Ansible can be used to automate much more than NSX-T related tasks.

With Ansible, you can also automate the build of the actual NSX-T core components, like the NSX-T Manager nodes and the (virtual) edge transport nodes. Ansible can also do the full day-0 deployment, including adding the proper license key, creating and applying the initial profiles, and creating a complete logical network infrastructure using Tier-0 gateways, Tier-1 gateways, and segments.

Terraform

Terraform is comparable to Ansible, in the sense that you can use a piece of code to deploy your NSX-T logical network infrastructure. Terraform uses “providers,” which are comparable to Ansible’s modules. They contain the code to perform CRUD operations using IaC. Terraform can only

perform operations on an NSX-T infrastructure that is already available. This means that the NSX-T Managers and the (virtual) edge transport nodes have to exist. Unlike Ansible, Terraform can't build the actual NSX-T core components.

The beauty of using Terraform or Ansible is that you can also roll back infrastructure builds in a matter of seconds. So, if you have used the code to deploy the infrastructure, the state is tracked and memorized, so you can destroy what you have just built with a single command.

PowerShell/PowerCLI

PowerCLI is Microsoft's CLI tool; it provides automated operations inside the Microsoft infrastructure. PowerCLI uses modules, which are comparable to Ansible's modules and Terraform's providers. They contain the code to perform CRUD operations using IaC. PowerCLI is a collection of PowerShell modules that can be used to perform CRUD operations on VMware products. NSX-T has its own module, which you can use to perform a limited set of automated tasks. The NSX-T PowerCLI module is not as extensive as the ones in Ansible or Terraform.

Summary

This chapter explained the VMware and non-VMware tools you have available to perform automation using NSX-T.

I would like to thank you for reading this book and I hope you have enjoyed it.

—Iwan Hoogendoorn

Index

A

- Ansible, [321, 322](#)
- Appliance Proxy Hub (APH),
[41, 50, 56, 214](#)
- Authentication header (AH), [159](#)

B

- Binding Information Base (BIB), [105](#)

C

- Central Control Plane (CCP),
[41, 50, 56, 64, 190, 294](#)
- Certificate authority (CA), [161](#)
- Certificate revocation list (CRL), [161](#)
- Create, Read, Update, and Delete
(CRUD) operations, [317](#)

D

- Data Plane Development Kit
(DPDK), [45, 190](#)
- Dead peer detection (DPD), [162](#)
- Destination NAT (DNAT), [89, 90](#)
- Distributed firewall (DFW), [7](#)
 - advanced settings, [15, 16](#)
 - architecture, [41, 42](#)

- autosave, [38, 39](#)
 - concepts, [9](#)
 - configuration, [12–14](#)
 - context-aware attributes, [29–31](#)
 - context profile, [27–29](#)
 - create group, [21](#)
 - enable/disable, [35](#)
 - ESXi data, [43, 44](#)
 - factors, [13](#)
 - features, [8](#)
 - groups, [20](#)
 - membership criteria, [22, 23](#)
 - rule actions, [33, 34](#)
 - rule parameters, [19, 20](#)
 - rules, [14, 19](#)
 - ruleset, [40, 41](#)
 - rule settings, [34, 35](#)
 - save/view, [35–37](#)
 - scope, [32, 33](#)
 - security policy categories, [11](#)
 - security policy overview, [10](#)
 - services, [23–27](#)
 - sources and destinations, [20](#)
 - time windows, [16–18](#)
- Distributed intrusion
 - detection, [53](#)
 - architecture, [56, 57](#)
 - configuration, [57, 58](#)

INDEX

- Distributed intrusion
 - detection (*cont.*)
 - requirements, 54
 - use case, 54
- Distributed router (DR), 97, 99
- DNAT rules, 92
- DNS service, 118
 - configuration, 123
 - forwarder, 119
 - forwarder benefits, 121
 - forwarder verification, 123
 - zones, 122
- Dynamic group membership, 22
- Dynamic Host Configuration Protocol (DHCP), 87, 108
 - architecture, 109, 110
 - relay, 116, 117
 - relay server, 118
 - segment, 114, 115
 - server profile, 111, 112
 - server status, 116
 - Tier-0/Tier-1 gateway, 112, 113
 - use cases, 110
 - workflow, 111

E

- East-west partner security, 72, 73
- Encapsulating security payload
 - header (ESP), 159
- Endpoint protection, 79
 - process, 79, 80
 - rule, 83
 - use case, 70, 79

F

- Federation
 - components, 249, 250, 252
 - consistency/operations, 248
 - consumption methods, 252, 253
 - definition, 247
 - GM configuration ownership, 254
 - LM configuration, 255
 - LM configuration ownership, 253
 - use cases, 247

G

- Gateway firewall, 44–46
 - architecture, 50, 51
 - policy, 47
 - predefined, 46, 47
 - rule configuration, 49
 - rule settings, 50
 - settings, 48
- Generic Routing Encapsulation (GRE), 157
- Global Manager (GM) federation
 - configuration
 - definition, 255
 - flow, 256
 - greenfield, 259–261
 - GUI location selector, 264
 - initial configuration
 - steps, 256–258
 - local management
 - cluster, 262, 263
 - logical topologies, 267–274, 276

- networking, tier-1
 - gateway, 265, 267
- RTEP, 276
- stretched layer-2 network, 277
- stretched layer-2 network
 - packetwalk, 277–285, 287
- stretched layer-2 network VNI
 - mapping, 277
- system overview, 265
- Guest introspection,
 - architecture, 83, 84

H

- High availability (HA), 164

I, J, K

- IDS events, 59, 60
- IDS profile, configuration, 55, 58
- IDS rules, configuration, 58, 59
- IDS signatures, 54–56
- Infrastructure as Code (IaC), 321
- Internet Engineering Task Force (IETF), 101
- Intrusion detection, 53, 56, 59, 69

L

- Layer-2 VPN (L2VPN)
 - NSX-T, 178
 - autonomous edge, 190
 - client configuration, 191
 - configuration steps, 182

- deployment
 - considerations, 179
 - hub/spoke, 179
 - packet flow, 180, 181
 - packet format, 180
 - peer compatibility matrix, 190
 - scalability, 181
 - segment configuration,
 - 187, 188, 193, 194
 - server configuration, 184
 - session configuration,
 - 185, 186, 192
 - supported clients, 189
- NSX-v
 - managed edge, 191
 - standalone edge, 191
 - requirements, 178
 - use cases, 178
- Layer-7 gateway firewall, 66
- Load balancing
 - algorithms, 153
 - application profile
 - configuration, 144
 - architecture, 128
 - component relation, 127, 128
 - configuration steps, 137
 - creation, 138, 139
 - deployment modes, 133–137
 - layer 4, 126
 - layer 4 virtual server creation, 141
 - layer 7, 126, 127
 - layer 7 virtual server
 - creation, 142, 143
 - monitors, 132, 133, 154–156

INDEX

Load balancing (*cont.*)
 persistence profile
 configuration, 145
 profile, 130, 131
 rule, 149, 150
 server pool configuration, 151
 server pools, 131, 132
 SNAT mode, 153, 154
 SSL modes, 146, 147
 SSL profile configuration,
 148, 149
 tier-1 gateway, 129
 use cases, 125
 virtual server creation, 139, 140
 virtual servers, 130

M

Micro-segmentation, 3, 4
 application level, 6
 approaches, 6
 benefits, 5
 functionality level, 7
 use cases, 5

N, O

NAT64
 limitations, 104
 packet flow, 105, 107
 requirements, 103, 104
 rules, 108
 tables, 105
 use cases, 103

NAT configuration, 94
NAT packet flow, 97
Network Address Translation
 (NAT), 86–88
Network introspection, 53, 70, 73
New virtual machine, 80, 81
North-south partner security, 71
NSX Intelligence
 alarms, 209
 appliance, 196
 definition, 195
 deployment, 198–202
 footprint, 197
 real network insight, 196
 recommendations, 205,
 207–209
 requirements, 197
 use cases, 196
 verification, 202
 visualization, 202, 203, 205
NSX-T
 alarms, 213, 215, 217, 219
 events, 213
 federation (*see* Federation)
 load balancer (*see* Load
 balancing)
 native monitoring tools, 212, 213
NSX-T federation security
 components, 288
 configuration steps, 294, 295
 global groups, 291
 GM firewall policy, 288, 289
 GM firewall rules, 289, 290
 GM group rules, 297, 298

- GM group span, [298, 299](#)
 - GM group types/Span, [295–297](#)
 - GM group, virtual machine tag, [299–301](#)
 - GM/LM section overlap, [290, 291](#)
 - regions, [292, 293](#)
 - tags, [293, 294](#)
 - use cases, [287](#)
 - NSX-T LDAP integration
 - authentication, [234](#)
 - benefits, [234](#)
 - identity source
 - configuration, [235](#)
 - LDAP, [233](#)
 - server configuration, [236, 238](#)
 - NSX-T network topology, [210–212](#)
 - NSX-T VPN services
 - IPsec
 - CA, [161](#)
 - configuration, [166](#)
 - deployment
 - considerations, [163](#)
 - DPD, [162](#)
 - HA, [164, 165](#)
 - IKE profile configuration, [170–172](#)
 - methods, [158](#)
 - modes, [160](#)
 - profile configuration, [173](#)
 - protocols/algorithms, [160](#)
 - scalability, [165](#)
 - service configuration, [167–169](#)
 - session configuration, [174–177](#)
 - types, [162](#)
 - use cases, [157, 158](#)
 - layer 3 IPsec, [157](#)
- ## P
- PowerShell/PowerCLI, [321, 322](#)
 - Pre-shared key (PSK), [161, 179](#)
 - Public cloud gateway (PCG), [309–311](#)
 - Public cloud integration, VMware hyperscaler, [303](#)
- ## Q
- Quarantine, [81, 82, 310](#)
- ## R
- Reflexive NAT, [91, 92, 95, 96](#)
 - Remote Tunnel Endpoints (RTEPs), [260, 266, 267](#)
 - Reputation score, [61](#)
 - Role-based access control (RBAC)
 - authentication policy setting
 - configuration, VIDM users
 - default roles, [243](#)
 - LDAP users, [245](#)
 - local users, [244](#)
 - policy settings, [241, 242](#)
 - vidm users, [244, 245](#)
 - local NSX-T users, [239, 240](#)
 - NSX-T user account types, [238](#)
 - policy management/access/authentication, [238](#)

S

- Security Associations (SA),
159, 171, 174
- security policy overview, 9, 10
- Service consumption, 76
 - redirection policies, 78, 79
 - service chain, 77, 78
 - service profile, 77
- Service deployments, 75, 76
- Service insertion, 69
- Service profile, endpoint
protection, 82
- Service registration, 74
- Service router (SR), 44, 45, 97,
99, 129
- Service virtual machines (SVMs), 71
- Source NAT (SNAT), 88, 89, 135, 153

T

- Terraform, 321, 322
- Traditional data center security
model, 1, 2

U

- URL analysis, 60
 - architecture, 63, 64
 - categories, 61
 - configuration, 65
 - context profile, 65, 66
 - dashboard, 67, 68
 - requirements, 61

- use case, 61, 62
- Webroot, 63

V

- Virtual interfaces (VIFs), 20, 22
- Virtual Network Interface Card
(vNIC), 111
- Virtual tunnel interface (VTI),
163, 176, 177
- VMware cloud management
platforms
 - Ansible, 321
 - cloud director, 319, 320
 - PowerShell/PowerCLI, 322
 - terraform, 321
 - vRA, 318
 - vRO, 319
- VMware hyperscaler
 - benefits, 303
 - definition, 303
 - deployment
 - footprint, 307, 308
 - NSX cloud
 - component, 309–311
 - deployment modes, 312, 313
 - deployment steps, 313–315
 - multi-cloud deployment,
311, 312
 - use cases, 309
 - providers, 305
 - use cases, 306
 - VM, 306

VMware Identity Manager (vIDM)
 configuration, [228–230](#)
 definition, [221](#)
 integration
 verification, [231](#)
 local login, [232, 233](#)
 login screen, [231, 232](#)
 NSX-T integration, [222, 224](#)
 NSX-T integration
 prerequisites, [223](#)
 OAuth client, [225, 226](#)
 services enabled, [231](#)

 SHA-256 certificate
 thumbprint, [227](#)
 vRealize Automation (vRA), [318](#)
 vRealize Orchestrator (vRO), [319](#)

W, X, Y

Webroot, [63, 65](#)

Z

Zero trust security model, [4–6, 51](#)